EXHIBIT 25

EXPERT REPORT OF CYNTHIA RUDIN, PH.D.

May 30, 2022

Flores, et al. v. Stanford, et al.,
No. 18-cv-02468 (VB) (JCM)

Table of Contents

I. <u>Scope of Services and Summary of Opinions</u>

1. I am a professor of computer science, electrical and computer engineering, statistical science, mathematics, and biostatistics & bioinformatics at Duke University. I have spent my career conducting research on machine learning tools that help humans make better decisions, and I have published extensively, including in the Journal of the Royal Statistical Society and the Harvard Data Science Review, on the intersection of statistics, criminal justice, and interpretable machine learning.

2. I have been retained by Cravath, Swaine & Moore LLP ("Cravath") on behalf of Plaintiffs Carlos Flores, Lawrence Bartley, Demetrius Bennett, Antonio Roman, Vintarra Martin, Terrance Anderson, and all others similarly situated ("Plaintiffs") to serve as an expert in *Flores, et al. v. Stanford, et al.*, No. 18-CV-02468 (VB) (JCM). The hourly rate for my time in this matter is $500 per hour. This compensation is not contingent on the outcome of this litigation or on the substance of the opinions in this report.

3. A list of the sources that I considered in preparing this report is attached as Appendix A.

4. I reserve the right to update my opinions if new materials become available during the course of this litigation. I further reserve the right to amend this report on instruction of counsel, in response to any deposition testimony, or as a result of any motion or court order that may impact the nature or scope of claims and issues in this litigation. I also reserve the right to respond to any opinions in Defendants' expert reports. If I am called upon to testify at hearing or trial, I also reserve the right to employ demonstrative exhibits that summarize facts or opinions that are disclosed in this report or new information that subsequently becomes available.

5. This report proceeds as follows. First, I provide an overview of the development of New York State's COMPAS Re-Entry instrument ("NY Re-Entry COMPAS"). Second, I set forth in detail the key principles of *appropriate model development*, including: (1) transparency (*e.g.*, calculations can be readily scrutinized and checked, and are not overly complex); (2) internal consistency (*e.g.*, inter-rater reliability); (3) limited use of subjective variables; (4) alignment with established criminological theory (*e.g.*, lower criminal history score should yield a lower risk score); (5) dependence on appropriate variables (*e.g.*, variables should *not* primarily correlate with race or class status or result in differential treatment on the basis of legal life choices); (6) selection of appropriate outcomes (*e.g.*, conviction rather than arrest); and (7) recognition of the possibility that people—especially juveniles—can change (*e.g.*, not severely penalizing youth). Third, I assess the extent to which NY Re-Entry COMPAS comports with these principles of model development. Fourth, I set forth in detail the key principles of *appropriate model validation*: the model should be (1) trained on the correct population; (2) validated out-of-sample; (3) validated for internal consistency; (4) trained and validated on sufficiently large sample sizes; (5) examined for fairness; and (6) monitored over time and updated as necessary. Fifth, I assess the extent to which NY Re-Entry COMPAS comports with these principles. Finally, I outline the attributes of a superior risk assessment scoring

model, for parole release decisions generally and for juvenile offenders in particular. The conclusions I reach in each of these sections are summarized immediately below:

6. **Opinion #1:** The risk and needs tool currently used for parole release decision-making in New York, NY Re-Entry COMPAS, does not comport with key principles of appropriate model development. It is proprietary and thus cannot be scrutinized, it does not allow for data entry and calculation errors to be identified and remedied, it uses variables and outcomes that are not appropriate for parole release decisions, it is overly complex, it appears to have serious errors in its formulation or documentation, and it severely penalizes youth, among other issues discussed in this report.

7. **Opinion #2:** NY Re-Entry COMPAS does not comport with key principles of appropriate model validation. It has never been validated out-of-sample, it was trained for a purpose different than the purpose for which it is used (and thus was trained on an unrepresentative population), and it has never been updated, among other issues discussed in this report.

8. **Opinion #3:** A superior model would be transparent, use appropriate variables and outcomes, be properly validated, be easy to calculate and use, and allow for data entry and calculation errors to be identified and corrected by users.

II. <u>Expert Qualifications</u>

9. I hold an undergraduate degree from the University at Buffalo in mathematical physics and music theory, and a PhD from Princeton University in applied and computational mathematics.

10. Prior to becoming a professor at Duke University, I held positions at the Massachusetts Institute of Technology ("MIT"), Columbia University, and New York University.

11. My work focuses on interpretable machine learning and its applications. I work on decision trees, sparse linear models and scoring systems, variable importance measures, causal inference methods, interpretable deep learning, dimension reduction, and methods that can incorporate domain-based constraints and other types of domain knowledge into machine learning. I apply these techniques to critical societal problems in criminology, healthcare, and energy grid reliability, as well as to materials science and computer vision. Many of my interpretable machine learning algorithms heavily rely on efficient discrete optimization techniques. Some of my major projects include:

- My team's work on optimal scoring systems (sparse linear models with integer coefficients) has been applied to many healthcare and criminal justice contexts. For example, our work on seizure prediction in ICU patients allows doctors to monitor 2.8 times more patients, helps to prevent severe brain damage in critically ill patients, and saved over $6 million at two major hospitals in 2018. This work won the 2019 INFORMS Innovative Applications in Analytics Award. My work on criminal justice scoring systems has shown that simple scoring systems can be as accurate for prediction of recidivism as more complicated machine learning methods if they are optimized carefully on appropriate datasets.

- I led a team in the first major effort to maintain an underground electrical distribution network using machine learning, in collaboration with Con Edison in New York City. This work won the 2013 INFORMS Innovative Applications in Analytics Award.

- My collaborators and I developed code for detecting crime series in cities. This methodology, called the Series Finder algorithm, was adapted by the New York Police Department ("NYPD"). The NYPD's application, called Patternizr, has been running live in New York City since 2016 and is used to determine whether each new crime is related to past crimes.

12. I am the recipient of the 2022 Squirrel AI Award for Artificial Intelligence for the Benefit of Humanity from the Association for the Advancement of Artificial Intelligence ("AAAI"). This is the most prestigious award in the field of artificial intelligence, similar to the Nobel Prize and the Turing Award, and it carries a $1 million prize. I am also a 2022 Guggenheim Fellow and a three-time winner of the INFORMS Innovative Applications in Analytics Award. I was named one of the "Top 40 Under 40" by Poets and Quants in 2015, and was named by Businessinsider.com as one of the 12 most impressive professors at MIT in 2015.

13. I am a fellow of the American Statistical Association, the Institute of Mathematical Statistics, and the Association for the Advancement of Artificial Intelligence ("AAAI"), and I am a past Chair of both the INFORMS Data Mining Section and the Statistical Learning and Data Science Section of the American Statistical Association.

14. I have served on committees for the National Institute of Justice, AAAI, the Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining ("ACM SIGKDD"), and the Defense Advanced Research Projects Agency ("DARPA"). I have served on three committees for the National Academies of Sciences, Engineering and Medicine, including the Committee on Applied and Theoretical Statistics, the Committee on Law and Justice, and the Committee on Analytic Research Foundations for the Next-Generation Electric Grid.

15. I have delivered keynote and invited speeches at several conferences including SIGKDD ("KDD"), the Society for Artificial Intelligence and Statistics ("AISTATS"), Conference on Digital Experimentation ("CODE"), Machine Learning in Healthcare ("MLHC"), Fairness, Accountability and Transparency in Machine Learning ("FAT-ML"), the European Conference on Machine Learning and Principles and Practices of Knowledge Discovery in Databases ("ECML-PKDD"), and the Nobel Conference. My work has also been featured in news outlets including the *New York Times*, *Washington Post*, *Wall Street Journal*, *Boston Globe*, *Businessweek*, and NPR.

16. I publish regularly in the *Journal of Machine Learning Research*, *Machine Learning Journal*, *Annals of Applied Statistics*, *AISTATS*, *NeurIPS*, *Operations Research*, *Management Science*, *Nature Machine Intelligence*, *Journal of the Royal Statistical Society*, *Annals of Statistics*, *Harvard Data Science Review*, and other top publications in

machine learning and statistics. My work has been cited over 11,000 times and my i-10 index is 106; this means that 106 of my papers have been cited more than 10 times.

17. My curriculum vitae is attached to this report as Appendix B, including a list of all publications authored in the previous 10 years. During the previous four years, I have not testified as an expert at trial or by deposition.

III. <u>Development of NY Re-Entry COMPAS</u>

18. In or around 2007, New York State ("NYS") enlisted Northpointe to develop a tool to determine the level of *parole supervision*. This tool was not designed for use in release decisions.[1]

19. NYS decided *not* to use the standard COMPAS tool ("COMPAS Core").[2] Rather, NYS requested that Northpointe develop new risk scales [redacted].[3] [redacted].[4]

20. [redacted] Northpointe developed completely new risk algorithms for a NYS re-entry risk tool ("NY Re-Entry COMPAS").[5] Northpointe purported to validate

[1] *See, e.g.*, Staley Dep. Tr. 19:2-18 (testifying that NY Re-Entry COMPAS was designed to "help assign parolees to caseloads based on the level of supervision that they required," and that Ms. Staley was not aware in 2012 that the instrument was also used in making parole release decisions); *id.* at 66:4-68:10 (testifying that "the primary purpose [of NY Re-Entry COMPAS] was for supervision decisions for people in the community on parole," particularly with respect to the allocation of staffing resources); *id.* at 240:3-9 (testifying that Ms. Staley was not aware of NY Re-Entry COMPAS being intended for any other use – beyond allocation of parole supervision resources – at the time it was developed); [redacted]

[2] *See, e.g.*, Staley Dep. Tr. 37:4-13 ("Q. But is it the case that New York State worked with Northpointe to make certain modifications to the reentry COMPAS? A. Yes. Q. And is it your understanding that the New York State reentry COMPAS is a unique instrument that's used only in New York State? A. That is my understanding.").

[3] *See, e.g.*, Staley Dep. Tr. 99:13-101:5 (noting that Northpointe developed new risk scales at the request of the New York State Department of Parole, [redacted]

[4] *See, e.g.*, Staley Dep. Tr. 68:22-69:10, 95:5-8; [redacted]

[5] *See supra* note 3; [redacted]

the risk algorithms [redacted].[6]

21. Northpointe shared the algorithms with NYS in multiple reports.[7] However, NYS lost track of these reports,[8] and for the past decade, no one at the New York State Department of Corrections and Community Supervision ("DOCCS") has known that these reports existed.[9] That is, for the past decade, nobody within NYS knew what the model was that generated risk and needs scores for the tens of thousands of persons to whom COMPAS was administered, and nobody understood what the inputs were or how the inputs were transformed into scores. This means that nobody could detect or remedy errors; no one within NYS knew the formulas for generating the scores, so there was no way of knowing what the score should be.

22. NYS initially started using NY Re-Entry COMPAS to inform parole supervision levels.[10]

---

[6] [redacted]

[7] *See, e.g.*, Staley Dep. Tr. 54:13-55:14 (discussing the fact that, although Northpointe shared the NY Re-Entry COMPAS algorithms with NYS more than ten years ago, NYS believed that the algorithms were proprietary to Northpointe and that NYS was not entitled to information about the algorithms); [redacted]

[8] *See id.*; [redacted]

[9] *See, e.g.*, Staley Dep. Tr. 49:11-17 ("Q. So is it fair to say that from the perspective of DOCCS, for the last ten years, the New York modified reentry COMPAS risk scores have been a black box? A. I would say that that is fair to say, generally speaking, yes. Yes."); *id.* 50:20-51:2 ("Q. [I]n terms of how to get from the answers to the questions on the COMPAS questionnaire to the risk score deciles, that process, so far as you know, has been a black box for DOCCS over the past ten years; right? A. Yes."); *id.* at 192:16-193:6 (agreeing that DOCCS was unaware that it had access to the NY Re-Entry COMPAS algorithms); *id.* at 195:23-196:15 (same).

[10] *See* Staley Dep. Tr. 70:10-71:5 (testifying that NY Re-Entry COMPAS was initially used to inform supervision levels, and that it was later used to inform parole release decisions); [redacted]

23. In 2011, the NY Executive Law was amended to require the Parole Board to use a risk assessment tool when making parole release decisions.[11] Also in 2011, NYS decided to use the existing NY Re-Entry COMPAS for this purpose, even though it was designed for determining supervision levels rather than making parole release decisions,[12] and despite believing that the algorithms were proprietary.[13]

24. Although the Parole Board has a regulation requiring the use of a "periodically-validated risk assessment instrument,"[14] [redacted][15] [redacted].[16]

IV. NY Re-Entry COMPAS Does Not Comport with Appropriate Principles of Model Transparency, Internal Consistency, and Dependence on Appropriate Variables in Criminal Justice Risk Scoring

i. Appropriate Principles

25. Following best practice, criminal justice risk scoring models should, at a minimum, have the following attributes:

a. Transparency

26. Criminal justice risk scoring models should be *transparent*, to those operating the tool, to those subject to the tool, and to the public.[17]

---

[11] *Compare* N.Y. Exec. L. § 259-c (eff. June 22, 2010) *with* N.Y. Exec. L. § 259-c (eff. March 31, 2011) (adding requirement that the Parole Board utilize a "risk and needs assessment instrument that would be administered to all inmates eligible for parole").

[12] *See, e.g.*, Staley Dep. Tr. 70:10-71:5 (testifying that Ms. Staley's understanding was that NY Re-Entry COMPAS "was initially used for studying supervision levels and that at a later point, following the legislation, it was then . . . included in the Parole Board package for consideration, among other factors.").

[13] *See, e.g.*, Staley Dep. Tr. 54:13-22 ("Q. Okay. Have you had a view over the last decade as to whether DOCCS was entitled to get information from Northpointe about how the risk score algorithms for the New York modified reentry COMPAS work? A. My understanding has been that it was proprietary information and that DOCCS or anyone else was not entitled to it.")

[14] 9 N.Y.C.R.R. § 8002.2(a).

[15] [redacted]

[16] [redacted]

[17] *See, e.g.*, Cynthia Rudin, et al., The Age of Secrecy and Unfairness in Recidivism Prediction, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.6ed64b30 [hereinafter "Age of Secrecy and Unfairness"]; Cynthia Rudin, et al., Broader Issues Surrounding Model Transparency in Criminal Justice Risk Scoring, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.038c43fe; *see also, e.g.*, Alex Chohlas-Wood, Understanding risk assessment instruments in criminal justice, Brookings Institute (June 19, 2020)

27. "Transparency" in this context means, at a minimum, that those administering the tools and the persons whose lives are affected by the tools have access to information regarding how the scores are generated, including the inputs and the methods for transforming those inputs into scores, such that they can meaningfully interpret and understand the tools and challenge any errors.

28. There is a consensus in the computer science field that, since these tools are used to help determine people's freedom, they should be subject to a high level of scrutiny.[18]

29. Transparency requires that the tool be simple enough that individuals scored by the tool are generally able to understand exactly how their scores were calculated and to check for the possibility of data errors. The individuals should also be able to challenge the applicability (or fairness) of the tool in their specific cases. Lack of transparency in risk scoring models for certain types of decisions may create due process issues,[19] and simpler models have tended to perform as well as or better than more complex models in practice because they are easier to use.[20] The formulas used in the tool should be straightforward to use and to check for unnatural dependences on variables, such as nonmonotonicity (*i.e.*, the function slopes up when it should slope down, or vice versa).[21]

("[A]ny algorithm used in a high-stakes policy context, such as criminal sentencing, should be transparent. This ensures that any interested party can understand exactly how a risk determination is made, a distinct advantage over human decision-making processes. In this way, transparency can help establish trust, and is an acknowledgement of the role these tools play in consequential, impactful decisions.")

[18] *See, e.g.*, Chohlas-Wood, *supra* note 17 ("[A]lgorithms, and the data used to generate their predictions, should be carefully examined for the potential that any group would be unfairly harmed by the outputs. Judges, prosecutors, and data scientists should critically examine each element of data provided to an algorithm—particularly the predicted outcomes—to understand if these data are biased against any community.")

[19] *See, e.g.*, Danielle Keats Citron & Frank Pasquale, The Scored Society: Due Process For Automated Predictions, 89 Wash. L. Rev. 1, 1 (2014) ("The American due process tradition should inform basic safeguards. Regulators should be able to test scoring systems to ensure their fairness and accuracy. Individuals should be granted meaningful opportunities to challenge adverse decisions based on scores miscategorizing them. Without such protections in place, systems could launder biased and arbitrary data into powerfully stigmatizing scores."); *see also* Rebecca Wexler, Code of Silence: How private companies hide flaws in the software that governments use to decide who goes to prison and who gets out, *Washington Monthly* (June 11, 2017), https://washingtonmonthly.com/2017/06/11/code-of-silence/ [hereinafter "Wexler 2017(a)"]; Rebecca Wexler, When a computer program keeps you in jail: How computers are harming criminal justice, *New York Times* (June 13, 2017), https://www nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice html [hereinafter "Wexler 2017(b)"]; Age of Secrecy and Unfairness, *supra* note 17.

[20] Nathan James, Congr. Rsch. Serv., R44087, Risk and Needs Assessment in the Criminal Justice System 10 (2015) ("However, there is substantial evidence available to suggest that relatively brief risk indices outperform longer, more complex models. For example, one study in Pennsylvania found that risk assessment accuracy was improved by using only 8 of the 54 factors in one commonly used instrument.") (internal quotations and alterations omitted).

[21] *See id; see also* Kris Henning & Ryan Labrecque, Introduction to Risk Assessment for Criminal Justice Administrators, presented at the Justice Reinvestment Summit (February 2017), https://pdxscholar.library.pdx.edu/cgi/viewcontent.cgi?article=1021&context=ccj_fac, at 42 ("How do you choose a risk scale for your agency? Ease of Use - risk scales that are complicated, costly, and time consuming are rarely adopted/sustained. Reliability - similar risk scores should be produced by different raters.")

30. Public access to the tool ensures that it can be scrutinized and evaluated by criminal justice researchers, who can compare the accuracy of different models, evaluate for various forms of bias, and investigate whether the model comports with legal and policy values.[22]

31. If the calculation is even slightly complicated (meaning it cannot be completed without a calculator), the piece of computer code used to generate the scores should also be public. This allows others to check that there are no discrepancies between the published formula and the formula that is actually used, since such discrepancies can occur. In particular, if the computer code disagrees with the published formula, it is not generally possible to determine how much a risk score depends on a particular variable. For instance, in connection with our paper *The Age of Secrecy and Unfairness*,[23] my colleagues and I questioned whether the mathematical formula for the COMPAS score in the Northpointe Practitioner's Guide to COMPAS Core (2012, 2015, 2019) was correct as written; we specifically questioned whether the formula was missing a nonlinear transformation of age. This is an example of a case where outsiders could not determine COMPAS's dependence on a key variable. Referencing its Practitioner's Guide, Northpointe responded: "Discussions of appropriate variable transformations are beyond its scope and would not further its goals; however, we note that the skewed age variable is an ideal candidate for a normalizing transformation."[24] Thus, Northpointe has admitted that it is possible that its mathematical formulas may be different on paper than in code. But, as shown here, we cannot have a debate about whether a particular transformation or correlation within the model is a good idea if we cannot know what the model is—this further underscores the importance of transparency.

32. Operators of the tool should understand its computations, to ensure that it is being used correctly and for the correct purpose, and to spot errors should any arise.[25]

---

[22] *See* Chohlas-Wood, *supra* note 17; *see also, e.g.*, Brandon Garrett, Justice in Forensic Algorithms, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.a6b843af ("What may be needed is . . . to require validation and analysis. Indeed, it would be far preferable if researchers . . . received the underlying validation data and implementation data for risk assessments, so private and independent researchers can assess the use of such algorithms. Further, defense access is critical, so that criminal defendants can meaningfully access evidence, understand it, and litigate it. We need open science for risk assessment.")

[23] *See* Age of Secrecy and Unfairness, *supra* note 17 (noting that it is possible for formulas in COMPAS documentation to disagree with COMPAS formulas used in practice).

[24] Eugenie Jackson & Christina Mendoza, Setting the Record Straight: What the COMPAS Core Risk and Need Assessment Is and Is Not, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.1b3dadaa.

[25] *See. e.g.*, Alexandra Chouldechova, Transparency and Simplicity in Criminal Risk Assessment, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.b9343eec. ("Meaningful algorithmic transparency can entail many things, including but not limited to: access to the training algorithm, the trained algorithm, the training data and its provenance, a justification of the prediction target, a clear articulation of the purpose for which the algorithm is designed and how it is intended to be used, and a reproducible validation study of the algorithm for the proposed use case. Each of these forms of transparency enables us to deliberate on the algorithm and its use in a given context. Having a simple model makes it easier to communicate the details of the

b. Internal consistency and non-proneness to human error

33. Criminal justice risk scoring models should be *internally consistent and not error prone*.

34. All variables should be checked for correctness when using the tool, in order to improve correctness of information and internal consistency. Two separate calculations of the tool on the same individual at the same point in time should yield identical results, so interrater reliability studies should be performed to assess internal consistency.[26, 27]

35. Errors in risk score calculations can have severe implications.[28]

---

trained algorithm, but it does not buy us much if those models predict the wrong outcome, are trained on nonrepresentative data, or are applied out of context.")

[26] *See* Why Inter-Rater Reliability Matters for Recidivism Risk Assessment, *Public Safety Risk Assessment Clearinghouse*, *Bureau of Justice Assistance*, https://bja.ojp.gov/sites/g/files/xyckuh186/files/media/document/pb-interrater-reliability.pdf [hereinafter "BJA Inter-Rater Reliability"]; Queensland Audit Office Report 14, Criminal justice system—reliability and integration of data (2016-2017), https://www.qao.qld.gov.au/sites/default/files/reports/integration_and_reliability_of_data.pdf, at 1 ("For criminal justice data to be useful, it must be reliable; otherwise it can lead to incorrect perceptions and decisions."); CSG Justice Center Staff, In Brief: Understanding Risk and Needs Assessment, *Justice Center—The Council of State Governments* (January 13, 2017), https://csgjusticecenter.org/2017/01/13/in-brief-understanding-risk-and-needs-assessment/ ("Assessments should be checked for inter-rater reliability (i.e., that two different staff members score the same individual the same way on the risk instrument) and intra-rater reliability (i.e., an individual staff member scores the same person the same way repeatedly.) High-quality assessments require well-trained staff to conduct the assessments, clear and periodically updated scoring guidelines, regular validation studies, and ongoing quality improvement exercises."); James Austin, The Proper and Improper Use of Risk Assessment in Corrections, *Federal Sentencing Reporter*, 16(3) (Feb. 2004), at 4 ("An Independent and Objective Researcher Must Conduct Inter-Reliability and Validity Tests.").

[27] Some studies have shown low inter-rater reliability for recidivism risk, while others have not. Thus, it is important to assess each system separately. *See, e.g.*, James Austin, et al., Reliability and validity study of the LSI-R risk assessment instrument, The Institute on Crime, Justice and Corrections (Jan 2003), https://www.ojp.gov/ncjrs/virtual-library/abstracts/reliability-and-validity-study-lsi-r-risk-assessment-instrument; Christopher Baird, A Question of Evidence: A Critique of Risk Assessment Models Used in the Justice System, *Special Report, National Council on Crime and Delinquency* (Feb. 2009); Michael Rocque & Judy Plummer-Beale, In the eye of the beholder? An examination of the inter-rater reliability of the LSI-R and YLS/CMI in a correctional agency, *Journal of Criminal Justice*, 42(6) (2014); Grant Duwe, The Development, Validity, and Reliability of the Minnesota Screening Tool Assessing Recidivism Risk (MnSTARR), *Criminal Justice Policy Review*, 25(5) (2014); Christopher Lowenkamp, et al., Assessing the Inter-rater Agreement of the Level of Service Inventory Revised, *Federal Probation*, 68(3); Leonard M. Van der Knaap, et al., Reevaluating Interrater Reliability in Offender Risk Assessment, *Crime & Delinquency*, 58(1) (2012).

[28] *See., e.g.*, Garrett, *supra* note 22 ("When these mistakes are made in entering data, because the COMPAS system is not transparent, a judge or a defendant cannot readily correct the error. No one is told that an error occurs, and if an error is brought to a judge's attention, a judge cannot know whether the error played a role in a risk recommendation. Further, in a pretrial setting in which decisions are made quickly, with limited information and often without meaningful defense representation, by the time an error is caught it may be too late to remedy the harm caused by an improper detention decision."); Greg Ridgeway, Transparency, Statistics, and Justice System Knowledge Is Essential for Science of Risk Assessment, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.cb0f8674 ("Bushway, Owens, and Piehl (2012) found that 10% of the worksheets used to extract sentencing recommendations for specific cases in Maryland had errors. These errors caused the calculated sentencing guidelines for prison sentences to be between one year too short to two years too long. Judges, unaware of these errors, imposed sentences that were shorter or longer depending on the direction of the

36. Tools that use manually-entered information are generally not as consistent as those using automatically-generated variables due to the high risk of typographical errors.[29]

c. **Non-subjectivity**

37. Criminal justice risk scoring models should *avoid questions eliciting subjective judgments or opinions*. This includes a broad range of subjective questions.

38. The tool should avoid non-objective questions answered by the person who administers the assessment. This includes questions requiring professional judgment[30] on topics such as [redacted][31] [redacted],[32] [redacted].[33]

39. Subjective questions can easily lead to a lack of internal consistency, as answers can change depending on arbitrary factors such as who is administering the questionnaire, including confusion regarding what the question is asking or even personal dislike of the subject of the COMPAS assessment.

40. Subjective questions may lead to or amplify racial or other biases on the part of the person who administers the assessment.[34]

41. Questions that require a future prediction [redacted] require humans to perform risk calculations mentally. Hence,

errors. Data quality matters when producing decision-support tools, but it also shows that such tools are powerful influencers of judges.")

[29] It is common knowledge that errors from data-entry can be extremely costly. *See, e.g.*, Brady Behrman, Automation Eliminates Expensive Data Entry Errors, *Data Science Central* (Sept. 23, 2020), https://www.datasciencecentral.com/automation-eliminates-expensive-data-entry-errors/ ("Even the best data entry professionals make significantly more errors than automated systems. Humans are prone to distraction and fatigue, and the most productive user cannot match the pace of the automatic transfer of data between integrated platforms . . . . Data entry *is* of strategic importance because faulty data is a significant risk.")

[30] *See* Austin, *supra* note 26, at 3 ("Unfortunately, it has also been shown that professional judgments are, by far, the least accurate risk assessment method. Too often, these judgments are no more than 'gut' reactions that often vary from expert to expert on the very same offender.")

[31] [redacted]

[32] [redacted]

[33] [redacted]

[34] *See, e.g.*, Taylor Walker, Psychological Evaluations And Other Subjective Assessments Contribute To Racial Disparities In Parole Decisions, Says Report, *Witness LA* (March 25, 2022), https://witnessla.com/psychological-evaluations-and-other-subjective-assessments-contribute-to-racial-disparities-in-parole-decisions-says-report/ ("After controlling for factors like the crime category, the individual's age, criminal history, and educational attainment, researchers found that prior 'professional assessments' accounted for nearly half of the disparity. Black people were more likely than white or Latino people to have their parole applications opposed by prosecutors, more likely to have disciplinary citations, and less likely to receive low-risk scores from psychiatrists.")

the calculation becomes subjective and inaccuracies in these predictions propagate to the risk score. The point of these instruments is to apply systematic, rationalized methods of predicting risk consistently to all cases. Using subjective opinion or judgment questions as the basis of a risk tool is contrary to that end.

42. Furthermore, it is not appropriate for individuals who admit to needing help to be penalized by being denied parole. Indeed, it is totally counterproductive.

d. Alignment with established criminological theory

43. Criminal justice risk scoring models should not disagree with well-studied criminological theory on the relationships between variables and outcomes.[35]

44. In particular, all else being equal, a longer criminal history should lead to a higher risk estimate, not a lower estimate.[36] This is called "monotonicity": estimates should increase (or decrease) as a function of specific variables.

e. Appropriate variables

45. Criminal justice risk scoring models should depend on appropriate variables.

46. Risk scores that assist in making decisions that impact bail, parole, or sentencing should not use subjective information, or information primarily designed to reveal class status, race, or gender. For instance, family support characteristics are not appropriate for risk scores,[37] but could inform other types of decisions (such as counseling needs).

47. Class status, race, or gender should not unfairly bias the score. For example, level of family support may not be under a person's control (*e.g.*, a neglectful parent is not the fault of the child), and inclusion of this factor subjects people to differential treatment on

---

[35] *See* Austin, *supra* note 26, at 4 (Table 5: "Examples of Risk Factors That Predict Recidivism").

[36] Risk scores generally combine points for various aspects of risk. Typically, increased criminal history leads to increased points, where the points contribute to the risk estimates. Increases in either juvenile crimes or total numbers of past crimes may increase risk estimates, but never decreases them. Typically, increased age leads to decreased risk estimates.

[37] Michael Tonry, Legal and Ethical Issues in the Prediction of Recidivism, *Federal Sentencing Review*, 26(3) (2014), https://scholarship.law.umn.edu/faculty_articles/525, at 167-69. There are at least two reasons to exclude "family support" variables: (1) possible racial bias, and (2) using these variables would force people to be treated differently "on the basis of law-abiding decisions about how to live their lives." Tonry discusses why such variables were excluded from the Salient Factor score: "The initial guidelines incorporated status variables such as employment, education, residential status, and family characteristics, but these were abandoned because they are heavily correlated with race. Blacks were on average less well educated than whites, had weaker employment records, had less stable residential circumstances, and had weaker family roles and statuses. Use of such factors would systematically treat blacks more severely than whites, and that, the U.S. Parole Commission decided, would be unjust." Tonry further explains: "Personal characteristics such as education, employment, residential stability, and family circumstances should not ordinarily be included among aggravating sentencing and parole criteria because their use systematically disadvantages black and other minority offenders (though they may be used as mitigating factors)." Tonry also lists several key principles for risk scores, such as "[d]on't treat people differently on the basis of law-abiding decisions about how to live their lives" and gives an example of family circumstance variables for this principle.

the basis of legal conduct, ordinary life choices, or a status that ought not be the source of further disadvantage.

f. Appropriate outcomes

48. Criminal justice risk scoring models should be developed using appropriate outcomes.

49. Outcome variables should generally be convictions rather than arrests, to avoid predicting police behavior or crimes that the subject is materially less likely to have committed.[38]

g. Recognition that people—especially juveniles—can change

50. Criminal justice risk scoring models should recognize the possibility that people can change. In particular, youth offenders may change as they mature into adulthood.[39]

51. Heavily weighting static variables such as age-at-first-arrest is particularly problematic because it does not recognize the possibility of change. Numerous studies show that recidivism rates decline with age,[40] whereas age-at-first-arrest does not change with age (indeed, by definition it can never change). Age-at-first arrest may also be correlated with socioeconomic status, as well as with race, given the correlation of police behavior with race. This is why some developers of risk scores have chosen not to use this factor.[41] In contrast to age-at-first-arrest, dynamic variables (*e.g.*, education and job

---

[38] *See. e.g.*, Sarah L. Desmarais, The Role of Risk Assessment in the Criminal Justice System: Moving Beyond a Return to the Status Quo, *Harvard Data Science Review*, 2(1) (2020), https://doi.org/10.1162/99608f92.181cd09f ("It is well documented that racial and ethnic minorities are more likely to be arrested for behaviors that would not result in arrest for those of other racial and ethnic backgrounds. As a result, many argue for criminal history to be operationalized using prior convictions."); [redacted]

[39] *See, e.g.*, Edward P. Mulvey, et al., An examination of change in dynamic risk of offending over time among serious juvenile offenders, *Journal of Criminal Justice*, 45 (2016), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4917019/, at 9 ("Several findings from these analyses demonstrate that there is potential value in considering shifts in risk indicators over time. First, it appears that risk scores for adolescents in the Pathways study changed over time. The scores generally depart more from the baseline with each successive interview, as would be expected if the scores are capturing dynamic risk and if actual changes on these indicators are occurring over time. Furthermore, overall risk scores decrease over time. The calculated risk scores at each follow-up interview generally decrease as time passes.")

[40] *See id.*; Matthew Clarke, Justice Department Releases Ten-Year Recidivism Study, *Prison Legal News*, March 1, 2022, at 50 ("As might be expected, those with the most prior arrests were the most likely to be arrested during the next ten years. But age was strongly correlated to arrest rate, with the youngest releasees being arrested at more than twice the rate of the oldest.")

[41] *See* Tonry, *supra* note 37 ("Age at first commitment was initially a factor but was later abandoned" due to concerns about racial bias); Paula J. Fite, et al., Explaining Discrepancies in Arrest Rates Between Black and White Male Juveniles, *Journal of Consulting and Clinical Psychology*, 77(5) (2009), at 920-21. In the latter study, the authors showed that race was a useful predictor of being arrested: "Black youths were significantly more likely to be arrested as juveniles than White youths across all types of arrests. The discrepancy was particularly pronounced for the variable representing any type of criminal charge as a juvenile. Specifically, approximately one half of Black boys in the sample had been charged with a crime as a juvenile, while approximately one third of White boys had a criminal charge as a juvenile." This shows that even if a factor is predictive in the statistical sense, there are reasons in justice to exclude it. The authors controlled for risk factors (*i.e.*, low academic achievement, poor parent–child

training) *do* change over time and can lead to the possibility of a lower risk score. Recognizing people's capacity for change, criminal justice risk scoring models should incorporate such variables.

ii. NY Re-Entry COMPAS Does Not Comport with Any of these Principles.

a. NY Re-Entry COMPAS is not transparent

52. NY Re-Entry COMPAS is a black box (*i.e.*, a tool lacking transparency), not only to parole applicants, but also to DOCCS itself; parole applicants do not have access to COMPAS proprietary formulas, and prior to this litigation, DOCCS personnel did not even know they had access to the formulas, much less receive training on how to understand those formulas.[42]

53. COMPAS is unusual in that, unlike most risk instruments, which are transparent, simple, and available to the public,[43] COMPAS is proprietary – there is no public access.[44] Thus, it is not possible for a wide range of scientists or the public to scrutinize it.

54. COMPAS's calculations are complicated. [redacted] It is unclear whether the computer code undergirding the Northpointe software is, in fact, directing the computer to input the correct variables, transform them accurately, and/or apply the appropriate coefficients. In fact, it is not currently possible to check whether the formulas have been programmed to agree with the proprietary Northpointe reports, even if one has access to the confidential formulas in those reports.

55. COMPAS reports are easily misunderstood. [redacted][45] [redacted]

communication, peer delinquency, neighborhood disadvantage, and neighborhood problems). They found that, "[a]s anticipated, race was weakly to strongly correlated with 10 of the 14 risk factors, suggesting that Black youths were more likely to display and experience the majority of risk factors than White youths." The authors further found that controlling for these risk factors, "[i]ndirect effects analysis indicated that [the risk factors] significantly accounted for the relation between race and violence-related arrest" and that race was no longer a useful factor in predicting recidivism when controlling for these risk factors.

[42] *See supra* note 9.

[43] *See., e.g.*, Garrett, *supra* note 22 ("There are several levels of transparency needed to assure adequate vetting of algorithms in criminal justice. Fortunately, COMPAS is an outlier in the area of risk assessment and decision makers seem to be electing to use more public and transparent (and free) options. Most crimogenic risk assessments in use are in fact quite simple, relying on largely static factors (like age and criminal history), and the underlying instruments are made available. In that sense, they may be transparent.")

[44] [redacted]

[45] The needs scales [redacted] were developed to capture needs during a parole supervision period, *see, e.g.*, Staley Dep. Tr. 151:10-14.

[redacted][46] *See* Figure 1, below.

[redacted]

Figure 1: [redacted][47]

56. Second, for each of the scales, the bar chart includes a number from one to ten. Particularly because these numbers correspond to qualitative descriptors such as "Low," "Medium" or "High" and "Unlikely," "Probable" or "Highly Probable,"[48] people making decisions on the basis of an individual's COMPAS scores could easily misinterpret these numbers to correspond to the *probability of reoffending*. They do not; they simply correspond to decile (*i.e.*, the individual's ranking relative to the individuals in the dataset used to create the model, when that dataset is divided into ten equal parts).[49]

b. NY Re-Entry COMPAS is not internally consistent and is error-prone

57. [redacted][50]

---

46 [redacted]

47 [redacted]

48 [redacted]

49 [redacted]

50 [redacted]

58. Furthermore, typographical errors in the inputs can have dramatic effects on COMPAS scores.[51] Obvious discrepancies between COMPAS score sheets and presentence reports have been observed, and such data issues appear to have led to less favorable results for some applicants.[52] DOCCS has expressed concern internally about this issue, but has never studied the extent to which typographical errors and inconsistent data entry influence risk scores.[53] COMPAS scores used in other states have also been shown to contain frequent typographical errors.[54]

c. NY Re-Entry COMPAS includes non-objective questions

59. Patently incorrect answers to discretionary (*i.e.*, subjective) questions on a NY Re-Entry COMPAS assessment can dramatically and unjustifiably increase one or more of a parole applicant's NY Re-Entry COMPAS scores. For example, [redacted][55] [redacted][56] [redacted][57] [redacted][58] Such errors are of course highly problematic.[59]

---

[51] *See* Age of Secrecy and Unfairness, *supra* note 17; Wexler 2017(a), *supra* note 19; Wexler 2017(b), *supra* note 19.

[52] [redacted]

[53] [redacted] *id.* at 233:3-8 (admitting that no one within DOCCS has been asked to study the error rate on NY Re-Entry COMPAS assessments).

[54] *See* Age of Secrecy and Unfairness, *supra* note 17; Wexler 2017(a), *supra* note 19; Wexler 2017(b), *supra* note 19.

[55] [redacted]

[56] [redacted]

[57] [redacted]

[58] [redacted]

[59] [redacted] *id.* at 233:3-8 (admitting that no one within DOCCS has been asked to study the error rate on NY Re-Entry COMPAS assessments).

60. [redacted][60]

d. NY Re-Entry COMPAS does not align with established criminological theory

61. The formulas used to create the NY Re-Entry COMPAS scores do not appear to agree with standard principles concerning how variables should contribute to the risk (*see* Section IV.i.d above), and appear to contain typographical errors. Having someone outside Northpointe sanity-check these formulas would have helped to avoid problems such as unnatural dependencies on criminal history, age, or other variables.

62. [redacted]

63. [redacted]

[redacted]

Figure 2: [redacted].[61]

[60] [redacted]

[61] [redacted]

64. [redacted][62]:

- [redacted]
- [redacted]
- [redacted]
- [redacted]

65. [redacted]

[redacted]

Figure 3: [redacted]

66. [redacted].[63] [redacted]

67. [redacted]

---

[62] [redacted]

[63] [redacted]

[redacted][64] [redacted][65]

68. [redacted][66] [redacted][67] I am unaware of any literature showing that persons who commit a misdemeanor assault and a homicide pose the same risk of future violence or recidivism.

69. In any event, it is clear that the Northpointe report setting forth the risk assessment formulas NYS has been using to inform parole release decisions for more than a decade was not constructed carefully. For example, [redacted][68] [redacted]

---

[64] [redacted]

[65] [redacted]

[66] [redacted]

[67] Marieke Liem, Homicide offender recidivism: A review of the literature, *Aggression and Violent Behavior*, 18(1) (2013), https://www.sciencedirect.com/science/article/pii/S1359178912000882 ("While there exists an abundance of research on the criminal histories of homicide offenders, little is known about their future criminal behavior.")

[68] [redacted]

e. NY Re-Entry COMPAS depends on inappropriate variables

70. NY Re-Entry COMPAS was developed for needs assessments and was not intended to be used for risk assessments.[69] As a result, it *depends on variables that are not appropriate* for parole decisions.

71. [redacted][70] [redacted][71] [redacted].[72]

f. NY Re-Entry COMPAS depends on inappropriate outcomes

72. [redacted][73] [redacted]

g. NY Re-Entry COMPAS does not sufficiently reflect the possibility that people—especially juveniles—can change

73. NY Re-Entry COMPAS severely penalizes youth. [redacted]

[69] *See supra* note 1.

[70] [redacted]

[71] [redacted]

[72] *See* Tonry, *supra* note 37.

[73] [redacted]

[redacted][74]

74. [redacted][75] [redacted][76] [redacted][77] [redacted]

75. [redacted]

76. [redacted]

77. [redacted][78] [redacted]

---

[74] [redacted]

[75] *See supra* note 74.

[76] [redacted]

[77] [redacted]

[78] [redacted]

[redacted]

V. The Validation and Assessment of NY Re-Entry COMPAS Does Not Comport with Appropriate Principles of Validation and Assessment for Criminal Justice Risk Scoring Models

i. Appropriate Principles

78. Following best practice,[79] at a minimum, the development and validation of criminal justice risk scoring models should have the following attributes:

a. Trained on the appropriate population

79. Using the appropriate population is critical to avoiding sample bias.[80] If the tool was developed on a sample from the wrong population, the sample may not sufficiently represent the correct population. For instance, risks calculated from individuals who have all lived part of their adults lives outside of prison do not necessarily generalize to individuals who have never lived as adults outside of prison.

80. Juvenile offenders are an important subgroup to consider because (1) they had not yet matured into adults prior to sentencing, (2) many of them may never have lived as adults outside of prison, providing them with a dramatically different transition to adulthood than other prisoners, and thus (3) their outcomes may more heavily reflect the quality of prison programs (*e.g.*, education, job training) than do those of other prisoners.

b. Validated out-of-sample

81. A criminal justice risk scoring model should be assessed using data *other than* the data from which it was created. This is extremely important, as it ensures that the tool generalizes beyond the sample used to create it in the first place.[81] In other words, external validation helps ensure that a tool has predictive power when used with data beyond just the original data on which the tool was originally trained.

c. Validated for internal consistency

82. In my experience, if evaluations of the same person at the same point in time have high variance, the tool is not suitable for use in practice.[82] This is especially true in the case of

---

[79] *See, e.g.*, Risk Validation: Public Safety Risk Assessment Clearinghouse, *Bureau of Justice Assistance*, https://bja.ojp.gov/program/psrac/validation/risk-validation (accessed April 20, 2022) [hereinafter "BJA Risk Validation"].

[80] *See* Sarah L. Desmarais & Jay P. Singh, Risk Assessment Instruments Validated and Implemented in Correctional Settings in the United States, *CSG Justice Center* (March 27, 2013), https://csgjusticecenter.org/wp-content/uploads/2020/02/Risk-Assessment-Instruments-Validated-and-Implemented-in-Correctional-Settings-in-the-United-States.pdf; Chouldechova, *supra* note 25.

[81] *See* BJA Risk Validation, *supra* note 79.

[82] *See* BJA Inter-Rater Reliability, *supra* note 26.

a re-entry tool, because decisions regarding an individual's freedom should not involve a high degree of randomness.

d. Trained and validated on sufficiently large sample sizes

83. Small training set sizes can lead to overfitting, and small validation sets may lead to extremely large (and thus less useful) confidence intervals. In particular, a sufficiently large sample of juvenile offenders is necessary to evaluate a risk score for juvenile offenders.

e. Examined for fairness

84. The tool should not predict significantly less accurately for any important subpopulation. Race and gender are important subpopulations, as are juvenile offenders.

f. Monitored over time and updated over time as necessary

85. Monitoring studies should consider not only predictive accuracy, but also internal consistency. That is, studies should consider how often the same person's record would be scored differently under random evaluation circumstances (*e.g.*, staff change, location in the state, etc.).[83]

ii. <u>NY Re-Entry COMPAS fails to comport with these principles.</u>

86. NY Re-Entry COMPAS was designed not to inform release decisions, but rather, to determine the level of parole supervision upon release. As outlined below, major negative consequences flow from this, [redacted]

a. [redacted]

87. [redacted][84] [redacted]

b. [redacted]

88. [redacted]

---

[83] *See* BJA Risk Validation, *supra* note 79 ("Agencies should plan to monitor and, if possible, improve the performance of tools on a periodic basis.")

[84] [redacted]

[redacted][85], [redacted][86]. [redacted][87] [redacted]

c. [redacted]

89. [redacted][88]

d. [redacted]

90. [redacted][89] [redacted]

---

[85] [redacted]

[86] [redacted]

[87] [redacted]

[88] *See supra* note 50.

[89] *See supra* notes 3-4.

[redacted][90]

91. [redacted][91] [redacted] Typical minimum sample sizes for logistic regression are at least 100 per class (*e.g.*, individuals arrested for violent offenses)[92] and at least 20 events (*e.g.*, violent offenses) per variable (EPV).[93] [redacted][94] [redacted][95]

92. [redacted][96] [redacted][97] [redacted]

---

[90] [redacted]

[91] [redacted]

[92] [redacted]

[93] Peter C. Austin & Ewout W. Steyerberg, Events per variable (EPV) and the relative performance of different strategies for estimating the out-of-sample validity of logistic regression models, *Statistical Methods in Medical Research*, 26(2) (2014) ("Differences between the bootstrap-corrected approach and the use of an independent validation sample were minimal once the number of events per variable was at least 20.")

[94] [redacted]

[95] [redacted]

[96] [redacted]

[97] [redacted]

[redacted][98]

e. [redacted]

93. [redacted][99] [redacted][100] [redacted][101] [redacted][102]

f. [redacted]

94. [redacted][103]

VI. Attributes of a Superior Criminal Justice Risk Scoring Model, for Parole Release Decisions Generally and for Juvenile Offenders in Particular

95. *Respectful of change:* As discussed above, the instrument should not be designed in such a way that current youthfulness or age at time of first arrest severely disadvantages the

---

[redacted]

[98] [redacted]

[99] [redacted]

[100] [redacted]

[101] [redacted]

[102] *See supra* note 98.

[103] *See, e.g., supra* note 15.

applicant. Similarly, the new instrument should not use formulas that penalize youth offenses in a permanent way (*i.e.*, in a way that is impossible to overcome). In recognition of the fact that people mature as they transition from adolescence to adulthood, risk scores should be constructed such that crimes committed as youths are less heavily weighted (*i.e.*, permitted to be outweighed by dynamic achievements) over time. This can be accomplished by using less steep functions of age, and by allowing dynamic risk-reducing accomplishments to factor more heavily for younger applicants. Such dynamic variables might include education and job training. [redacted].[104] After a fixed number of years have elapsed since an applicant has reached the age of maturity, the new instrument should have little impact from: juvenile arrests, juvenile convictions, Juvenile Offender, Youthful Offender, or family court adjudications; any cases sealed pursuant to CPL §§ 160.50 and 160.55; and any cases that did not terminate in a criminal conviction.

96. *Transparency*: The new instrument should not be proprietary, and all inputs, variables, coefficients, formulas, algorithms, code, and scoring rules should be freely, readily, and publicly available. This will ensure that parole applicants and their advocates have the opportunity to understand how risk scores are derived, to detect and correct errors, to take immediate and meaningful steps to improve such scores and thereby to expedite, refine, and intensify rehabilitative efforts. Also, the formulas should not be overly complicated: logarithms, exponents, and complicated interaction terms should be avoided if possible. Such complicated terms are more difficult to use: they make it harder to check whether a score has been computed correctly for a particular individual;[105] they make it difficult to check whether the formula has been correctly input into the computer; and they make it more difficult to check whether the variables have been used appropriately (*i.e.*, if the functions are monotonic and align with criminological theory, as discussed above). Complex terms also make it more difficult to check whether the relative weighting of terms (*e.g.*, age versus dynamic factors) is appropriate, and make it more difficult to teach the formulas to the people and organizations who administer the tool (*e.g.*, prison personnel), who interpret the scores (*e.g.*, parole boards; prosecutors), and whose liberty depends in part on those scores (*e.g.*, parole applicants).

97. *Development at scale*: The new instrument should be constructed using an adequately large dataset from the appropriate population, and this dataset should oversample juvenile offenders to ensure an adequately large sample size with respect to this subpopulation.

98. *Outcomes*: For purposes of identifying unfavorable parole outcomes within the dataset, the risk scores should be developed based on convictions rather than arrests, as an individual's arrest for a particular crime does not sufficiently indicate that the individual actually committed that crime or the charge for which he or she was arrested. As

---

[104] *See* Tonry, *supra* note 37; Fite, *supra* note 41.

[105] Age of Secrecy and Unfairness, *supra* note 17; Wexler 2017(a), *supra* note 19; Wexler 2017(b), *supra* note 19. As discussed above, typographical errors do appear to occur in COMPAS score calculations.

discussed above, designing predictive models from arrests leads implicitly to predicting police behavior, which could lead to higher dependence on variables correlated with race.

99. *Validation population*: The instrument should be validated on a dataset other than the dataset used to create it. This dataset should be adequately large and should oversample juvenile offenders to ensure an adequately large sample size with respect to this subpopulation. The validation sample should be from the correct population.

100. *Periodic validation*: The new instrument should be re-validated on a comparable dataset at least once every five years.

101. *Errors, non-objective, and inappropriate questions*: The new instrument's means of deriving a parole applicant's raw scores (*e.g.*, questionnaire; self-assessment; input of criminal and disciplinary history) should not rely on non-objective questions about the following topics: [redacted];[106] [redacted];[107] [redacted];[108] [redacted][109] [redacted][110] In addition, DOCCS staff should be prohibited from changing the answers given by incarcerated persons under any self-reporting section.[111] The new instrument should generally avoid subjective questions. Furthermore, in order to minimize human error, the new instrument should not employ manual input of criminal records and disciplinary data, and should instead import such data automatically from the New York State Division of Criminal Justice Services.

102. *Inter-rater reliability studies*: Inter-rater reliability studies should be conducted in connection with the new instrument. Inter-rater reliability should be high, so that an individual's score will be consistently calculated no matter who is administering the assessment and/or computing the score.

103. *Error correction:* There should be an easy process by which an error in a score can be corrected prior to a decision being made on the basis of that score.

104. *Racial and class bias*: The new instrument's design should be non-discriminatory on the basis of race and class status.

105. *Probability estimates rather than decile scores:* In order to avoid the possibility that someone might interpret a decile score as a probability (*e.g.*, interpret a decile 7 score as

---

106 [redacted]

107 [redacted]

108 [redacted]

109 [redacted]

110 [redacted]

111 [redacted]

a 70% probability of recidivism), the new tool should present the scores in a manner that does not create confusion between deciles and probability.

VII. Conclusion

106. In this report, I have discussed both appropriate principles for the design of a criminal justice risk scoring model, as well as appropriate validation and testing for such a model. Although these principles have been extensively studied, NY Re-Entry COMPAS violates them in obvious ways. NY Re-Entry COMPAS's key violations of these principles are as follows:

- It is proprietary, making it extremely difficult to scrutinize. Errors in the data used to generate the model, errors in the model formulas, errors in the code implementing the model, and/or errors in the input into the model when calculating a specific person's scores can easily propagate without anyone being able to detect them.
- It depends quite heavily on age in a way that does not recognize the possibility that people can mature and change. As a result, the instrument severely and inappropriately penalizes younger people.
- [redacted] This directly contravenes standard criminological theory.
- It uses subjective information. This is problematic for reasons of inter-rater reliability issues and racial bias. No study has been performed to examine this.
- It has never undergone a validation study using a separate test dataset. This means that it is impossible to know how well the instrument performs on new data.

107. Adopting an instrument that possesses the attributes set forth in Section VI above will help avoid these and other issues moving forward.

Respectfully Submitted,

Cynthia Rudin
May 30, 2022

Appendix A:
Materials Considered

Case Materials
Complaint, *Flores et al. v. Stanford et al.,* No. 7:18 Civ. 02468 (VB) (JCM) (S. D. N. Y.)

Discovery Documents
Northpointe 001 - Northpointe 008 Northpointe 00112 - Northpointe 00182 STANFORD0056884 - STANFORD0056926 STANFORD0086845 - STANFORD0086898 STANFORD0068673 - STANFORD0068673 STANFORD0068674 - STANFORD0068771 STANFORD0068772 - STANFORD0068772 STANFORD0068773 - STANFORD0068776 STANFORD0068777 - STANFORD0068780 [redacted] Copy of ReEntry ParoleRisk RawScoresGivenNYSIDs 11032020 (003).xlsx

Deposition Transcripts and Exhibits
Transcription of Deposition of Michele Staley – DOCCS 30(b)(6) (December 2, 2021) and Exhibits
Transcription of Deposition of Jennifer Bryant (December 16, 2021) and Exhibits
[redacted]

Statutes and Regulations
N.Y. Exec. L. § 259-c (eff. June 22, 2010)
N.Y. Exec. L. § 259-c (eff. March 31, 2011)
9 N.Y.C.R.R. § 8002.2(a)

Publicly Available Documents
Alex Chohlas-Wood, Understanding risk assessment instruments in criminal justice, Brookings Institute (June 19, 2020)
Alexandra Chouldechova, Transparency and Simplicity in Criminal Risk Assessment, Harvard Data Science Review, 2(1) (2020)
Brady Behrman, Automation Eliminates Expensive Data Entry Errors, Data Science Central (Sept. 23, 2020)
Brandon Garrett, Justice in Forensic Algorithms, Harvard Data Science Review, 2(1) (2020)
CSG Justice Center Staff, In Brief: Understanding Risk and Needs Assessment, Justice Center—The Council of State Governments (January 13, 2017)

Publicly Available Documents
Christopher Baird, A Question of Evidence: A Critique of Risk Assessment Models Used in the Justice System, Special Report, National Council on Crime and Delinquency (Feb. 2009)
Christopher Lowenkamp, et al., Assessing the Inter-rater Agreement of the Level of Service Inventory Revised, Federal Probation, 68(3) (2012)
Cynthia Rudin, et al., The Age of Secrecy and Unfairness in Recidivism Prediction, Harvard Data Science Review, 2(1) (2020)
Cynthia Rudin, et al., Broader Issues Surrounding Model Transparency in Criminal Justice Risk Scoring, Harvard Data Science Review, 2(1) (2020)
Danielle Keats Citron & Frank Pasquale, The Scored Society: Due Process For Automated Predictions, 89 Wash. L. Rev. 1, 1 (2014)
Edward P. Mulvey, et al., An examination of change in dynamic risk of offending over time among serious juvenile offenders, Journal of Criminal Justice, 45 (2016)
Eugenie Jackson & Christina Mendoza, Setting the Record Straight: What the COMPAS Core Risk and Need Assessment Is and Is Not, Harvard Data Science Review, 2(1) (2020)
Grant Duwe, The Development, Validity, and Reliability of the Minnesota Screening Tool Assessing Recidivism Risk (MnSTARR), Criminal Justice Policy Review, 25(5) (2014)
Greg Ridgeway, Transparency, Statistics, and Justice System Knowledge Is Essential for Science of Risk Assessment, Harvard Data Science Review, 2(1) (2020)
James Austin, et al., Reliability and validity study of the LSI-R risk assessment instrument, The Institute on Crime, Justice and Corrections (Jan 2003)
James Austin, The Proper and Improper Use of Risk Assessment in Corrections, Federal Sentencing Reporter, 16(3) (Feb. 2004)
Julia Angwin, et al., Machine Bias, ProPublica (May 23, 2016)
Kris Henning & Ryan Labrecque, Introduction to Risk Assessment for Criminal Justice Administrators, presented at the Justice Reinvestment Summit (February 2017)
Leonard M. Van der Knaap, et al., Reevaluating Interrater Reliability in Offender Risk Assessment, Crime & Delinquency, 58(1) (2012)
Marieke Liem, Homicide offender recidivism: A review of the literature, Aggression and Violent Behavior, 18(1) (2013)
Matthew Clarke, Justice Department Releases Ten-Year Recidivism Study, Prison Legal News, March 1, 2022
Michael Rocque & Judy Plummer-Beale, In the eye of the beholder? An examination of the inter-rater reliability of the LSI-R and YLS/CMI in a correctional agency, Journal of Criminal Justice, 42(6) (2014)
Michael Tonry, Legal and Ethical Issues in the Prediction of Recidivism, Federal Sentencing Review, 26(3) (2014)
Nathan James, Congr. Rsch. Serv., R44087, Risk and Needs Assessment in the Criminal Justice System 10 (2015)
Paula J. Fite, et al., Explaining Discrepancies in Arrest Rates Between Black and White Male Juveniles, Journal of Consulting and Clinical Psychology, 77(5) (2009)

Publicly Available Documents
Peter C. Austin & Ewout W. Steyerberg, Events per variable (EPV) and the relative performance of different strategies for estimating the out-of-sample validity of logistic regression models, Statistical Methods in Medical Research, 26(2) (2014)
Prisoner Hunger Strike Solidarity Coalition, Advice for prisoners and their supporters regarding Board of Parole Hearings psychological evaluations, San Francisco Bay View (April 7, 2017)
Queensland Audit Office Report 14, Criminal justice system— reliability and integration of data (2016-2017)
Rebecca Wexler, Code of Silence: How private companies hide flaws in the software that governments use to decide who goes to prison and who gets out, Washington Monthly (June 11, 2017)
Rebecca Wexler, When a computer program keeps you in jail: How computers are harming criminal justice, New York Times (June 13, 2017)
Risk Validation: Public Safety Risk Assessment Clearinghouse, Bureau of Justice Assistance, (accessed April 20, 2022)
Sarah L. Desmarais & Jay P. Singh, Risk Assessment Instruments Validated and Implemented in Correctional Settings in the United States, CSG Justice Center (March 27, 2013)
Sarah L. Desmarais, The Role of Risk Assessment in the Criminal Justice System: Moving Beyond a Return to the Status Quo, Harvard Data Science Review, 2(1) (2020)
Shawn D. Bushway, Emily G. Owens & Anne Morrison Piehl, Sentencing Guidelines and Judicial Discretion: Quasi-experimental Evidence from Human Calculation Errors, (2012)
Taylor Walker, Psychological Evaluations And Other Subjective Assessments Contribute To Racial Disparities In Parole Decisions, Says Report, Witness LA (March 25, 2022)
Why Inter-Rater Reliability Matters for Recidivism Risk Assessment, Public Safety Risk Assessment Clearinghouse, Bureau of Justice Assistance

Northpointe Data Files
Absconder Warrants.por
All Warrants.por
Arrests Match.por
NorthPointe Sample.por
SRP Match.por

All additional materials cited or referenced in my report.

Appendix B:
Curriculum Vitae

Cynthia Rudin

Professor of Computer Science, Electrical and Computer Engineering,
Statistical Science, Mathematics, and Biostatistics & Bioinformatics
Duke University
D342 LSRC Building, Box 90129, Durham, NC 27708
cynthia@cs.duke.edu https://users.cs.duke.edu/ cynthia/

Research Interests

My work focuses on interpretable machine learning, which is crucial for responsible and trustworthy AI. This includes the design of algorithms for interpretable modeling, interpretable policy design, variable importance measures, causal inference methods, uncertainty quantification, and methods that can incorporate domain-based constraints and other types of domain knowledge into machine learning. These techniques are applied to critical societal problems in healthcare, criminal justice, energy grid reliability, and other areas. The interpretable machine learning algorithms heavily rely on efficient discrete optimization techniques.

I am the winner of the 2022 *Squirrel AI Award for Artificial Intelligence for the Benefit of Humanity* from the Association for the Advancement of Artificial Intelligence (AAAI). This is the most prestigious award in artificial intelligence. Similar only to world-renowned recognitions, such as the Nobel Prize and the Turing Award, it carries a monetary reward at the million-dollar level. I am also a 2022 Guggenheim fellow. Some accomplishments:

- My collaborators and I developed the first practical code for decision lists and decision trees that provably optimize accuracy and sparsity. This work had an oral presentation at KDD 2017, a spotlight at NeurIPS 2019, and recent work appeared in ICML 2020 and AAAI 2022.
- Our work on optimal scoring systems (sparse linear models with integer coefficients) has been applied to several important healthcare and criminal justice applications. For instance, my collaboration's work on seizure prediction in ICU patients allowed for better allocation of ICU patient monitoring resources. This work won the 2019 INFORMS Innovative Applications in Analytics Award. In a validation study, it allowed doctors to monitor 2.8X more patients and saved over $6M at two major hospitals in FY2018. It helps to prevent severe brain damage in critically ill patients.
- I led a team in the first major effort to maintain an underground electrical distribution network using machine learning, in joint work with Con Edison in NYC. This won the 2013 INFORMS Innovative Applications in Analytics Award.
- I solved a well-known (previously open) theoretical problem in machine learning as a PhD student, which is whether AdaBoost maximizes the margin like SVM. Subsequent work solved a published COLT open problem.
- My collaborators and I developed code for detecting crime series in cities. This methodology (specifically, the Series Finder algorithm) was adapted by the NYPD and their application (Patternizr) has been running live in NYC since 2016 for determining whether each new crime is related to past crimes.
- I enjoy competing in data science competitions and coaching student teams. We have won awards in several competitions, including the FICO Recognition Award for the first Explainable Machine Learning Competition in 2018, NTIRE Superresolution competition in 2018, PoeTix Literary Turing Competition in 2018, and we placed second in one of the ICML Exploration and Exploitation competitions.
- I have given invited and keynote talks at KDD (2014 and 2019), AISTATS, ECML-PKDD, ML in Healthcare, FAT-ML (Fairness, Accountability, and Transparency), and the Nobel Conference. My work frequently appears in the media, including the NY Times, Washington Post, Boston Globe, Wall Street Journal, MSNBC, and National Public Radio.

Education

- **Ph.D.: Princeton University**. Program in Applied and Computational Mathematics. Title: **Boosting, Margins, and Dynamics**. Advisors: Ingrid Daubechies and Robert Schapire.
- **BS/BA: University at Buffalo (SUNY), Honors Program**. Outstanding Senior Award in the Arts and Sciences (one awarded per year university-wide), separate outstanding senior awards from the Physics Department, Mathematics Department, and Music Department, BS Mathematical Physics, BA Music Theory, Minor in Computer Science, Summa Cum Laude, 1999

Employment History

- **Duke University**, Computer Science Department (50%), Electrical and Computer Engineering Department (50%), Secondary appointments in Statistical Science, Biostatistics & Bioinformatics, and Mathematics. Associate Professor 2016-2019, Professor 2019-present
- **Massachusetts Institute of Technology**, MIT Computer Science and Artificial Intelligence Laboratory and Sloan School of Management, Associate Professor of Statistics 2013-2016, Assistant Professor of Statistics 2009-2013.
- **Columbia University**, Center for Computational Learning Systems, Associate Research Scientist, 2007-2009.
- **NSF Postdoctoral Research Fellow, New York University**, 2004-2007

Peer-Reviewed Publications

Papers associated with major awards (winner and finalist)

1. Gah-Yi Ban and Cynthia Rudin. *MSOM Best OM paper in OR Award, INFORMS. (Awarded to the best paper in Operations Research within the last 3 years, awarded by Manufacturing and Service Operations Management Society of INFORMS.)* **The Big Data Newsvendor: Practical Insights from Machine Learning**, Operations Research, Vol. 67, No. 1, pages 90-108, 2019.
2. Aaron F. Struck, Berk Ustun, Andres Rodriguez Ruiz, Jong Woo Lee, Suzette LaRoche, Lawrence J. Hirsch, Emily J Gilmore, Jan Vlachy, Hiba Arif Haider, Cynthia Rudin, M Brandon Westover. *2019 INFORMS Innovative Applications in Analytics Award* (shared with paper below). **Association of an Electroencephalography-Based Risk Score With Seizure Probability in Hospitalized Patients**, JAMA Neurology, 74 (12), 1419-1424, 2017.
3. Berk Ustun and Cynthia Rudin. *2019 INFORMS Innovative Applications in Analytics Award*, also *2017 INFORMS Computing Society Student Paper Prize.* **Learning Optimized Risk Scores**. Journal of Machine Learning Research, 2019. Shorter version **Learning Optimized Risk Scores from Large-Scale Datasets**. Knowledge Discovery in Databases (KDD), 2017.
4. Chaofan Chen, Kangcheng Lin, Cynthia Rudin, Yaron Shaposhnik, Sijia Wang, Tong Wang. **A Holistic Approach to Interpretability in Financial Lending: Models, Visualizations, and Summary-Explanations.** Decision Support Systems, In press, 2021.
   Preliminary work:
   - Chaofan Chen, Kangcheng Lin, Cynthia Rudin, Yaron Shaposhnik, Sijia Wang, Tong Wang. *Winner of the FICO Recognition Award for the Explainable Machine Learning Challenge, 2018.* **An Interpretable Model with Globally Consistent Explanations for Credit Risk.** NIPS 2018 Workshop on Challenges and Opportunities for AI in Financial Services: the Impact of Fairness, Explainability, Accuracy, and Privacy, 2018.
5. Cynthia Rudin and Berk Ustun. *Finalist for 2017 Daniel H. Wagner Prize for Excellence in Operations Research, Institute for Operations Research and Management Science (INFORMS).* **Optimized Scoring Systems: Towards Trust in Machine Learning for Healthcare and Criminal Justice**. INFORMS Journal on Applied Analytics, Special Issue: 2017 Daniel H. Wagner Prize for Excellence in Operations Research Practice, 48(5), pages 399–486, September-October, 2018.
6. William Souillard-Mandar, Randall Davis, Cynthia Rudin, Rhoda Au, David J. Libon, Rodney Swenson, Catherine C. Price, Melissa Lamar, Dana L. Penney. *2016 INFORMS Innovative Applications in Analytics Award* (shared with paper below). **Learning Classification Models of Cognitive Conditions from Subtle Behaviors in the Digital Clock Drawing Test**. Machine Learning, volume 102, number 3, 2016.
7. Berk Ustun and Cynthia Rudin. *2016 INFORMS Innovative Applications in Analytics Award* and also *Runner up, Invenia Labs SEE Award 2018 - Supporting Machine Learning Research with a Positive Impact on Social, Economic, or Environmental (SEE) Challenges*. **Supersparse Linear Integer Models for Optimized Medical Scoring Systems**. Machine Learning, volume 102, number 3, 2016.
8. Cynthia Rudin, Şeyda Ertekin, Rebecca Passonneau, Axinia Radeva, Ashish Tomar, Boyi Xie, Stanley Lewis, Mark Riddle, Debbie Pangsrivinij, Tyler McCormick. *2013 INFORMS Innovative Applications in Analytics Award* **Analytics for Power Grid Distribution Reliability in New York City**. INFORMS Journal on Applied Analytics, volume 44, issue 4, pages 364-383, 2014.

9. Indraneel Mukherjee, Cynthia Rudin, and Robert E. Schapire. *This paper answered an open question published in COLT 2010.* **The Rate of Convergence of AdaBoost**, Journal of Machine Learning Research, volume 14, pages 2315-2347, August 2013.
   Preliminary work:
   - Indraneel Mukherjee, Cynthia Rudin, and Robert E. Schapire. **The Rate of Convergence of AdaBoost**, Proceedings of the 24th Annual Conference on Learning Theory (COLT), 2011.

Publications and other best paper awards (not including those above)

2022

10. Lesia Semenova, Cynthia Rudin, and Ron Parr. **On the Existence of Simpler Machine Learning Models.** ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT), 2022.
11. Jiachang Liu, Chudi Zhong, Margo Seltzer, and Cynthia Rudin. **Fast Sparse Classification for Generalized Linear and Additive Models**, AISTATS, 2022.
12. Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, Chudi Zhong. **Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges**, Statistics Surveys, 2022.
13. Hayden McTavish, Chudi Zhong, Reto Achermann, Ilias Karimalis, Jacques Chen, Cynthia Rudin, Margo Seltzer. **How Smart Guessing Strategies Can Yield Massive Scalability Improvements for Sparse Decision Tree Optimization**, AAAI, 2022.
14. Zhi Chen, Alexander Ogren, Chiara Daraio, L. Catherine Brinson, Cynthia Rudin. *Winner of the 2022 Physical and Engineering Sciences (SPES) and the Quality and Productivity (Q&P) Student Paper Competition of the American Statistical Association.* **How to See Hidden Patterns in Metamaterials with Interpretable Machine Learning**.
15. Tong Wang and Cynthia Rudin. **Subgroup Identification for Enhanced Treatment Effect with Decision Rules**, INFORMS Journal on Computing, 2022.
16. Yaron Shaposhik and Cynthia Rudin. **Globally-Consistent Rule-Based Summary-Explanations for Machine Learning Models: Application to Credit-Risk Evaluation**. Accepted with minor revision to the Journal of Machine Learning Research.
    - An earlier version appeared at the Conference on Information Systems and Technology (CIST), 2019.
17. Caroline Wang, Bin Han, Bhrij Patel, Feroze Mohideen, Cynthia Rudin. **In Pursuit of Interpretable, Fair and Accurate Machine Learning for Criminal Recidivism Prediction**. Journal of Quantitative Criminology, 2022.
18. Masoud Afnan, Michael Anis Mihdi Afnan, Yanhe Liu, Julian Savulescu, Abhishek Mishra, Vincent Conitzer, Cynthia Rudin. **Data solidarity for machine learning for embryo selection; a call for the creation of an open access repository of embryo data**. Reproductive BioMedicine Online, 2022.
19. Chunxiao Li, Cynthia Rudin, and Tyler McCormick. **Rethinking Nonlinear Instrumental Variable Models through Prediction Validity**, Journal of Machine Learning Research, 2022.

2021

20. Cynthia Rudin and Shawn Bushway. **A Truth Serum for your Personal Perspective on Facial Recognition Software in Law Enforcement**, Translational Criminology, Fall 2021
21. Tianyu Wang and Cynthia Rudin. *Best Paper Finalist for INFORMS Data Mining Best Paper Competition - General Track 2021*, **Bandit Learning for Proportionally Fair Allocations**, 2021.
22. Yingfan Wang, Haiyang Huang, Cynthia Rudin, and Yaron Shaposhnik. **Understanding How Dimension Reduction Tools Work: An Empirical Approach to Deciphering t-SNE, UMAP, TriMAP, and PaCMAP for Data Visualization**, Journal of Machine Learning Research, 2021.
23. Tianyu Wang, Marco Morucci, M. Usaid Awan, Yameng Liu, Sudeepa Roy, Cynthia Rudin, Alexander Volfovsky. **FLAME: A Fast Large-scale Almost Matching Exactly Approach to Causal Inference**. Journal of Machine Learning Research, 2021.

- Software for FLAME by our students Vittorio Orlandi and Neha Gupta was the *honorable mention for the 2022 John M. Chambers Statistical Software Award from the American Statistical Association.*

24. Stefano Tracà, Cynthia Rudin, Weiyu Yan. *Best paper award, INFORMS 2016 Data Mining & Decision Analytics Workshop,* also *Finalist for 2015 IBM Service Science Best Paper Award.* **Regulating Greed over Time in Multi-Armed Bandits.** Journal of Machine Learning Research, 2021.
25. Jianyou Wang, Xiaoxuan Zhang, Yuren Zhou, Christopher Suh, Cynthia Rudin. **There Once Was a Really Bad Poet, It Was Automated But You Didn't Know It.** Transactions of the Association for Computational Linguistics (TACL), 2021. (Presented at ACL, 2021)
26. Divya Koyyalagunta, Anna Sun, Rachel Lea Draelos, Cynthia Rudin. **Playing Codenames with Language Graphs and Word Embeddings.** Journal of Artificial Intelligence Research (JAIR), 2021.
27. Michael Anis Mihdi Afnan, Cynthia Rudin, Vincent Conitzer, Julian Savulescu, Abhishek Mishra, Yanhe Liu and Masoud Afnan. **Ethical Implementation of Artificial Intelligence to Select Embryos in In Vitro Fertilization.** Fourth AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2021. (Our related work "Embryo selection using Artificial Intelligence (AI): Epistemic and ethical considerations" accepted for oral presentation at European Society of Human Reproduction and Embryology, ESHRE, 2021, and our related work "Embryo selection by "black-box" artificial intelligence: The ethical and epistemic considerations" accepted for oral presentation at Fertility Society of Australia and New Zealand Annual Conference, 2021)
28. Michael Anis Mihdi Afnan, Yanhe Liu, Vincent Conitzer, Cynthia Rudin, Abhishek Mishra, Julian Savulescu, Masoud Afnan. **Interpretable, Not Black-Box, Artificial Intelligence Should be Used for Embryo Selection**, Human Reproduction Open, 2021.
29. Alina Jade Barnett, Fides Regina Schwartz, Chaofan Tao, Chaofan Chen, Yinhao Ren, Joseph Y. Lo, Cynthia Rudin. **IAIA-BL: A Case-based Interpretable Deep Learning Model for Classification of Mass Lesions in Digital Mammography**, Nature Machine Intelligence, 2021.
30. Alina Jade Barnett, Fides Regina Schwartz, Chaofan Tao, Chaofan Chen, Yinhao Ren, Joseph Y. Lo, and Cynthia Rudin. **Interpretable Mammographic Image Classification using Cased-Based Reasoning and Deep Learning**, Oral Presentation, IJCAI-21 Workshop on Deep Learning, Case-Based Reasoning, and AutoML: Present and Future Synergies, 2021.
31. Zhicheng Guo, Cheng Ding, Xiao Hu, and Cynthia Rudin. **A Supervised Machine Learning Semantic Segmentation Approach for Detecting Artifacts in Plethysmography Signals from Wearables**. Physiological Measurement, 2021.
32. Alex Oesterling, Angikar Ghosal, Haoyang Yu, Rui Xin, Yasa Baig, Lesia Semenova, and Cynthia Rudin. **Duke Data Science 3C Citation Classification Task B**, (Outlines our strategy for our *third-place competition entry in the 2021 3C Shared Task Citation Context Classification based on Purpose competition*), Oral Presentation, Second Workshop on Scholarly Document Processing (SDP) Workshop, NAACL, 2021

2020

33. Zhi Chen, Yijie Bei and Cynthia Rudin. **Concept Whitening for Interpretable Image Recognition.** Nature Machine Intelligence, 2020.
34. Jiayun Dong and Cynthia Rudin. **Variable Importance Clouds: A Way to Explore Variable Importance for the Set of Good Models**, Nature Machine Intelligence, 2020
35. Beau Coker, Cynthia Rudin and Gary King. **A Theory of Statistical Inference for Ensuring the Robustness of Scientific Results.** Management Science, 2020.
36. Jimmy Lin, Chudi Zhong, Diane Hu, Cynthia Rudin, Margo Seltzer. **Generalized and Scalable Optimal Sparse Decision Trees.** ICML, 2020.
37. Tianyu Wang and Cynthia Rudin. **Bandits for BMO Functions.** ICML, 2020.
38. Marco Morucci, Vittorio Orlandi, Sudeepa Roy, Cynthia Rudin, Alexander Volfovsky. **Adaptive Hyper-box Matching for Interpretable Individualized Treatment Effect Estimation.** UAI, 2020.

39. Cynthia Rudin, Caroline Wang and Beau Coker. **The *Age* of Secrecy and Unfairness in Recidivism Prediction.** Harvard Data Science Review, 2020.

    - Broader Issues Surrounding Model Transparency in Criminal Justice Risk Scoring. (Rejoinder) Harvard Data Science Review, 2020.

40. Sachit Menon, Alexandru Damian, Nikhil Ravi, Shijia Hu, Cynthia Rudin. **PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models**. CVPR, 2020.
41. Marco Morucci, Vittorio Orlandi, Sudeepa Roy, Alexander Volfovsky. **Almost-Matching-Exactly for Treatment Effect Estimation under Network Interference**. AISTATS, 2020.
42. Tianyu Wang, Weicheng Ye, Dawei Geng and Cynthia Rudin. **Towards Practical Lipschitz Bandits**. Foundations of Data Science (FODS), 2020.
43. Jerry Liu, Nathan O'Hara, Alexander Rubin, Rachel Draelos, Cynthia Rudin. **Metaphor Detection Using Contextual Word Embeddings from Transformers**. Proceedings of the Second Workshop on Figurative Language Processing, pages 250-255, 2020.
44. Hunter Gregory, Steven Li, Pouya Mohammadi, Natalie Tarn, Rachel Draelos, Cynthia Rudin. **A Transformer Approach to Contextual Sarcasm Detection in Twitter.** Proceedings of the Second Workshop on Figurative Language Processing, pages 270-275, 2020.
45. Qinwen Huang, Ye Zhou, Xiaochen Du, Reed Chen, Jianyou Wang, Cynthia Rudin, Alberto Bartesaghi. **Cryo-ZSSR: Multiple-Image Super-Resolution Based on Deep Internal Learning**. Machine Learning for Structural Biology Workshop at the 34th Conference on Neural Information Processing Systems, 2020.

2019

46. Cynthia Rudin and Joanna Radin. **Why Are We Using Black Box Models in AI When We Don't Need To? A Lesson From an Explainable AI Competition**. Harvard Data Science Review, 2019.
47. Aaron Fisher, Cynthia Rudin, Francesca Dominici. **All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously**. Journal of Machine Learning Research, 2019.
48. Harsh Parikh, Cynthia Rudin, and Alexander Volfovsky. **An Application of Matching After Learning To Stretch (MALTS) to the ACIC 2018 Causal Inference Challenge Data**. Observational Studies, Issue 5, pages 118-130, 2019.
49. Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, Cynthia Rudin. **This Looks Like That: Deep Learning for Interpretable Image Recognition**. NeurIPS spotlight (top 3% of papers), 2019.
50. Xiyang Hu, Cynthia Rudin, and Margo Seltzer. **Optimal Sparse Decision Trees**. NeurIPS spotlight (top 3% of papers), 2019.
51. Fernanda Bravo, Cynthia Rudin, Yuting Yuan and Yaron Shaposhnik. **Simple Rules for Predicting Congestion Risk in Queueing Systems: Application to ICUs**, 2019 INFORMS Workshop on Data Science (DS 2019), oral presentation.
52. Peter Hase, Chaofan Chen, Oscar Li, Cynthia Rudin. **Interpretable Image Recognition with Hierarchical Prototypes**. AAAI Human Computation (AAAI-HCOMP), 2019.
53. Cynthia Rudin. **Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead.** Nature Machine Intelligence, 2019.

    - Shorter version called **Please Stop Explaining Black Box Machine Learning Models for High Stakes Decisions** appeared at NIPS 2018 Workshop on Critiquing and Correcting Trends in Machine Learning, 2018.

54. Stefano Tracà, Weiyu Yan, and Cynthia Rudin. **Reducing exploration of dying arms in mortal bandits**. UAI, 2019
55. Usaid Awan, Yameng Liu, Marco Morucci, Sudeepa Roy, Cynthia Rudin, and Alexander Volfovsky. **Interpretable Almost-Exact Matching With Instrumental Variables.** UAI, 2019

56. Awa Dieng, Yameng Liu, Sudeepa Roy, Cynthia Rudin, and Alexander Volfovsky. **Interpretable Almost-Exact Matching For Causal Inference.** AISTATS, 2019.

2018

57. Fulton Wang, Tyler McCormick, Cynthia Rudin, and John Gore. *Best Poster Award, Conference of the ASA Section on Statistical Learning and Data Mining, 2014.* **Modeling Recovery Curves With Application to Prostatectomy.** Biostatistics, 2018.
58. Cynthia Rudin and Yining Wang. *Finalist for 2017 QSR (Quality, Reliability and Statistics) best refereed paper competition, INFORMS 2017.* **On Direct Learning to Rank and Rerank**. Artificial Intelligence and Statistics (AISTATS), 2018.
59. John Benhart, Tianlin Duan, Peter Hase, Liuyi Zhu, Cynthia Rudin. *Winner of the 2018 PoetiX Literary Turing Test Award for computer-generated poetry.* **Shall I Compare Thee to a Machine-Written Sonnet? An Approach to Algorithmic Sonnet Generation,** 2018.
60. Yijie Bei, Alex Damian, Shijia Hu, Sachit Menon, Nikhil Ravi, and Cynthia Rudin. *NTIRE-CVPR 2018 Image Super-Resolution Challenge: winner for Track 1 (classic bicubic), honorable mention for Track 2 (realistic mild adverse conditions).* **New Techniques for Preserving Global Structure and Denoising with Low Information Loss in Single-Image Super-Resolution**, New Trends in Image Restoration and Enhancement Workshop and Challenges on Super-Resolution, Dehazing, and Spectral Reconstruction, NTIRE-CVPR, 2018.
61. Cynthia Rudin and Şeyda Ertekin. **Learning Customized and Optimized Lists of Rules with Mathematical Programming**. Mathematical Programming C (Computation), Mathematical Programming Computation, Volume 10, Number 4, pages 659-702, 2018.
62. Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. **Learning Certifiably Optimal Rule Lists for Categorical Data**, Journal of Machine Learning Research, volume 18, no. 234, pages 1-78, 2018.
63. Nicholas Larus-Stone, Elaine Angelino, Daniel Alabi, Margo Seltzer, Vassilios Kaxiras, Aditya Saligrama and Cynthia Rudin. **Systems Optimizations for Learning Certifiably Optimal Rule Lists**. Conference on Machine Learning and Systems (MLSys), 2018.
64. Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. **Deep Learning for Case-based Reasoning through Prototypes: A Neural Network that Explains its Predictions**. Association for the Advancement of Artificial Intelligence (AAAI), 2018.
65. Chaofan Chen and Cynthia Rudin. **An Optimization Approach to Learning Falling Rule Lists**. Artificial Intelligence and Statistics (AISTATS), 2018.

2017

66. Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. **Certifiably Optimal Rule Lists for Categorical Data**, Knowledge Discovery in Databases (KDD - oral presentation), 2017.
67. Berk Ustun, Lenard A. Adler, Cynthia Rudin, Stephen V. Faraone, Thomas J. Spencer, Patricia Berglund, Michael J. Gruber, Ronald C. Kessler. **The World Health Organization Adult Attention-Deficit/Hyperactivity Disorder Self-Report Screening Scale for DSM-5**. JAMA Psychiatry, April 2017.
68. Tong Wang, Cynthia Rudin, Finale Doshi, Yimin Liu, Erica Klampfl, and Perry MacNeille. **Bayesian Rule Sets for Interpretable Classification, with Application to Context-Aware Recommender Systems**. Journal of Machine Learning Research (JMLR), volume 18, number 70, pages 1-37, 2017.
69. Fulton Wang and Cynthia Rudin. **Causal Falling Rule Lists**, Fairness, Accountability, and Transparency (FATML), 2017 (longer version on ArXiv).
70. Hongyu Yang, Cynthia Rudin, and Margo Seltzer. *Winner of Student Paper Competition sponsored by the Statistical Learning and Data Mining section (SLDM) of the American Statistical Association, 2016.* **Scalable Bayesian Rule Lists**. International Conference on Machine Learning (ICML), 2017.
71. Himabindu Lakkaraju and Cynthia Rudin. *Finalist for 2017 INFORMS Data Mining Best Paper Competition.* **Learning Cost-Effective and Interpretable Treatment Regimes**. Artificial Intelligence and Statistics (AISTATS), 2017.

- Shorter versions accepted to workshops: NIPS Machine Learning for Healthcare (ML4HC), NIPS Workshop on Interpretable Machine Learning in Complex Systems, and NIPS Workshop on ML and the Law, NIPS 2016.

2016

72. Jiaming Zeng, Berk Ustun, and Cynthia Rudin. *Winner of 2015 Undergraduate Statistics Research Project Competition (USRESP) sponsored by the American Statistical Association (ASA) and the Consortium for Advancement of Undergraduate Statistics Education (CAUSE).* **Interpretable Classification Models for Recidivism Prediction.** Journal of the Royal Statistical Society Series A, volume 180, issue 3, 2016.
73. Tong Wang, Cynthia Rudin, Finale Doshi, Yimin Liu, Erica Klampfl, and Perry MacNeille. **Bayesian Rule Sets for Interpretable Classification**. IEEE International Conference on Data Mining (ICDM), 2016.
74. Ramin Moghaddass, Cynthia Rudin, and David Madigan. **The Factorized Self-Controlled Case Series Method: An Approach for Estimating the Effects of Many Drugs on Many Outcomes.** Journal of Machine Learning Research, 17(185):1-24, 2016.
75. Berk Ustun, M. Brandon Westover, Cynthia Rudin, and Matt T. Bianchi. **Clinical Prediction Models for Sleep Apnea: The Importance of Medical History over Symptoms**. Journal of Clinical Sleep Medicine, volume 12, number 2, 2016.
76. Vikas Garg, Cynthia Rudin, and Tommi Jaakola. **CRAFT: ClusteR-specific Assorted Feature selecTion**, Artificial Intelligence and Statistics (AISTATS), 2016.
77. Benjamin Letham, Portia A. Letham, Cynthia Rudin, Edward P. Browne. **Prediction Uncertainty and Optimal Experimental Design for Learning Dynamical Systems.** Chaos, volume 26, number 6, 2016.
78. Benjamin Letham, Lydia M. Letham and Cynthia Rudin. **Bayesian Inference of Arrival Rate and Substitution Behavior from Sales Transaction Data with Stockouts**. Knowledge Discovery in Databases (KDD), 2016
79. Edward P. Browne, Benjamin Letham, and Cynthia Rudin. **A Computational Model of Inhibition of HIV-1 by Interferon-Alpha**. PLoS ONE, vol 11, no 3, pages 1–16, March, 2016.
80. Şeyda Ertekin and Cynthia Rudin. **A Bayesian Approach to Learning Scoring Systems**. Big Data, volume 3, number 4, 2016.

2015

81. Fulton Wang and Cynthia Rudin. *Winner of Best Student Paper Competition, Statistical Learning and Data Mining section (SLDM) of the American Statistical Association, 2015,* also *Finalist for INFORMS Data Mining Section Best Student Paper Award, 2015.* **Falling Rule Lists**. Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS), 2015.
82. Benjamin Letham, Cynthia Rudin, Tyler H. McCormick, and David Madigan. *Winner of Data Mining Best Student Paper Competition, INFORMS 2013,* also *Winner of Student Paper Competition sponsored by the Statistical Learning and Data Mining section (SLDM) of the American Statistical Association, 2014.* **Building Interpretable Classifiers with Rules using Bayesian Analysis: Building a Better Stroke Prediction Model**. Annals of Applied Statistics, volume 9, number 3, pages 1350-1371, 2015.
    Shorter version:
    - Benjamin Letham, Cynthia Rudin, Tyler McCormick, and David Madigan. **An Interpretable Model for Stroke Prediction Using Rules and Bayesian Analysis.** Proceedings of 2014 KDD Workshop on Data Science for Social Good, 2014
83. Tong Wang, Cynthia Rudin, Daniel Wagner and Rich Sevieri. *Second place in INFORMS 2015 Doing Good with Good OR Paper Competition.* **Finding Patterns with a Rotten Core: Data Mining for Crime Series with Core Sets**. Big Data. volume 3, issue 1, 2015. (Special issue on Data for Social Good)
84. Ramin Moghaddass and Cynthia Rudin. **The Latent State Hazard Model, with Application to Wind Turbine Reliability**. Annals of Applied Statistics, volume 9, number 4, pages 1823–1863, 2015.

85. Şeyda Ertekin, Cynthia Rudin, and Tyler McCormick. **Reactive Point Processes: A New Approach to Predicting Power Failures in Underground Electrical Systems**. Annals of Applied Statistics, volume 9, No 1, pages 122–144, 2015.

2014

86. Theja Tulabandhula and Cynthia Rudin. **Tire Changes, Fresh Air and Yellow Flags: Challenges in Predictive Analytics for Professional Racing**. Big Data, vol 2 issue 2, pages 97-112, June 20, 2014.
87. Şeyda Ertekin, Cynthia Rudin and Haym Hirsh. **Approximating the Crowd**. Data Mining and Knowledge Discovery, volume 28 issue 5-6, pages 1189-1221, September 2014.
    - Shorter versions appeared at NIPS Workshop on Computational Social Science and the Wisdom of Crowds, 2011, Proceedings of Collective Intelligence (CI) 2012, and Proceedings of the 2012 AAAI Fall Symposium on Machine Aggregation of Human Judgment, MAGG-2012.
88. Been Kim, Cynthia Rudin, and Julie Shah. **The Bayesian Case Model: A Generative Approach for Case-Based Reasoning and Prototype Classification**. Neural Information Processing Systems (NIPS), 2014.
89. Theja Tulabandhula, Cynthia Rudin. **On Combining Machine Learning with Decision Making**. Machine Learning (ECML-PKDD journal track), volume 93, Pages 33-64, 2014
90. Siong Thye Goh and Cynthia Rudin. **Box Drawings for Learning with Imbalanced Data**. Proceedings of 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2014.
91. Theja Tulabandhula and Cynthia Rudin. **Generalization Bounds for Learning with Linear, Polygonal, Quadratic and Conic Side Knowledge**. Machine Learning (ECML-PKDD journal track), December, 2014, pages 1-34.
    - Shorter version: Theja Tulabandhula, Cynthia Rudin. **Generalization Bounds for Learning with Linear and Quadratic Side Knowledge**. Proceedings of ISAIM 2014.
92. Jonathan Huggins and Cynthia Rudin. **Towards a Theory of Pattern Discovery** Proceedings of SIAM Conference on Data Mining (SDM) 2014.
93. Been Kim and Cynthia Rudin. **Learning About Meetings**, Data Mining and Knowledge Discovery, (ECML-PKDD Journal track), volume 28 issue 5-6, pages 1134-1157, September 2014.

2013

94. Benjamin Letham, Cynthia Rudin and Katherine Heller. **Growing a List**. Data Mining and Knowledge Discovery (DAMI), ECML-PKDD journal track. volume 27, pages 372-395, 2013.
95. Tong Wang, Cynthia Rudin, Daniel Wagner, Richard Sevieri. **Learning to Detect Patterns of Crime**, Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD), 2013.
    - Ideas from this work were implemented by the NYPD by Alex Chohlas-Wood and E.S. Levine in their algorithm Patternizr, which operates live in New York City.
96. Cynthia Rudin, Benjamin Letham, and David Madigan. **Learning Theory Analysis for Association Rules and Sequential Event Prediction**. Journal of Machine Learning Research (JMLR), volume 14, pages 3385-3436, 2013.
    - Shorter version: Cynthia Rudin, Ben Letham, Ansaf Salleb-Aouissi, Eugene Kogan, and David Madigan. **Sequential Event Prediction with Association Rules**, Proceedings of the 24th Annual Conference on Learning Theory (COLT), 2011.
97. Benjamin Letham, Cynthia Rudin and David Madigan. **Sequential Event Prediction**. Machine Learning, volume 93, pages 357-380, 2013
98. Theja Tulabandhula and Cynthia Rudin. *Finalist, Data Mining Best Student Paper Competition, INFORMS 2012.* **Machine Learning with Operational Costs**. Journal of Machine Learning Research (JMLR), volume 14, pages 1989-2028, July 2013. Preliminary work is in the following conference paper.

- Theja Tulabandhula and Cynthia Rudin. **The Influence of Operational Costs on Estimation**, Proceedings of the International Symposium on Artificial Intelligence and Mathematics (ISAIM), 2012.

2012

99. Tyler McCormick, Cynthia Rudin, and David Madigan. **Hierarchical Models for Association Rule Mining: A New Approach for Adverse Event Prediction in Clinical Trials**, Annals of Applied Statistics, volume 6, No. 2, pages 652–668, 2012.
100. Cynthia Rudin, David Waltz, Roger N. Anderson, Albert Boulanger, Ansaf Salleb-Aouissi, Maggie Chow, Haimonti Dutta, Philip Gross, Bert Huang, Steve Ierome, Delfine Isaac, Arthur Kressner, Rebecca J. Passonneau, Axinia Radeva, Leon Wu. **Machine Learning for the New York City Power Grid**, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 34, No 2, February 2012. *Spotlight Paper for the February 2012 Issue.*
101. Allison Chang, Cynthia Rudin, Mike Cavaretta, Robert Thomas and Gloria Chou. **Reverse-Engineering Quality Ratings**, Machine Learning: volume 88, issue 3, pages 369-398, 2012.

2011

102. Cynthia Rudin, Rebecca J. Passonneau, Axinia Radeva, Steve Ierome, and Delfina Isaac. **21st-Century Data Miners Meet 19th-Century Electrical Cables**, IEEE Computer, volume 44 no. 6, pages 103-105, June 2011. *(One of three articles featured on the cover of the magazine.)*
103. Şeyda Ertekin and Cynthia Rudin. **On Equivalence Relationships Between Classification and Ranking Algorithms**, Journal of Machine Learning Research, volume 12, pages 2905–2929, 2011.

2010

104. Cynthia Rudin, Rebecca J. Passonneau, Axinia Radeva, Haimonti Dutta, Steve Ierome, and Delfina Isaac. **A Process for Predicting Manhole Events in Manhattan**. Machine Learning, volume 80, pages 1–31, 2010.
    - Also oral presentation at ICML 2012

The following conference papers are also related to my projects on grid reliability.

- Rebecca J. Passonneau, Cynthia Rudin, Axinia Radeva, Ashish Tomar, Boyi Xie. **Treatment Effect of Repairs to an Electrical Grid: Leveraging a Machine Learned Model of Structure Vulnerability**, Proceedings of the KDD Applications in Sustainability (SustKDD) Workshop on Data Mining, 17th Annual ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2011.
- Dingquan Wang, Rebecca Passonneau, Michael Collins and Cynthia Rudin. **Modeling Weather Impact on a Secondary Electrical Grid**, 4th International Conference on Sustainable Energy Information Technology (SEIT-2014), 2014.
- Leon Wu, Timothy Teräväinen, Gail Kaiser, Roger Anderson, Albert Boulanger, and Cynthia Rudin. **Estimation of System Reliability Using a Semiparametric Model**, Proceedings of IEEE EnergyTech, 2011.
- Leon Wu, Gail Kaiser, Cynthia Rudin, and Roger Anderson. **Data Quality Assurance and Performance Measurement of Data Mining for Preventive Maintenance of Power Grid**, Proceedings of the KDD Workshop on Data Mining for Service and Maintenance (KDD4Service), 17th Annual ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2011.
- Leon Wu, Gail Kaiser, Cynthia Rudin, David Waltz, Roger Anderson, Albert Boulanger, Ansaf Salleb-Aouissi, Haimonti Dutta, and Manoj Poolery. **Evaluating Machine Learning for Improving Power Grid Reliability**, Proceedings of the ICML 2011 workshop on Machine Learning for Global Challenges, International Conference on Machine Learning, 2011.
- Axinia Radeva, Cynthia Rudin, Rebecca Passonneau and Delfina Isaac. **Report Cards for Manholes**, Proceedings of the International Conference on Machine Learning and Applications (ICMLA), 2009. *Best Poster Award.*

- Rebecca Passonneau, Cynthia Rudin, Axinia Radeva and Zhi An Liu. **Reducing Noise in Labels and Features for a Real World Dataset: Application of NLP Corpus Annotation Methods**, Proceedings of the 10th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing), 2009.
- Haimonti Dutta, Cynthia Rudin, Becky Passonneau, Fred Seibel, Nandini Bhardwaj, Axinia Radeva, Zhi An Liu, Steve Ierome, Delfina Isaac. **Visualization of Manhole and Precursor-Type Events for the Manhattan Electrical Distribution System**, Workshop on GeoVisualization of Dynamics, Movement and Change, 11th AGILE International Conference on Geographic Information Science, 2008.
- Boyi Xie, Rebecca J. Passonneau, Haimonti Dutta, Jing-Yeu Miaw, Axinia Radeva, Ashish Tomar, Cynthia Rudin. **Progressive Clustering with Learned Seeds: An Event Categorization System for Power Grid.** 24th International Conference on Software Engineering and Knowledge Engineering (SEKE 2012). Redwood City, CA. July 1-3, 2012.

2009

105. Cynthia Rudin. **The P-Norm Push: A Simple Convex Ranking Algorithm that Concentrates at the Top of the List**, Journal of Machine Learning Research, volume 10, pages 2233–2271, 2009.

  - Shorter version: Cynthia Rudin. **Ranking with a P-Norm Push**. Proceedings of the Nineteenth Annual Conference on Learning Theory (COLT), pages 589 - 604, 2006.

  An application of the P-Norm Push is described in this conference paper:

  - Heng Ji, Cynthia Rudin, and Ralph Grishman. **Re-ranking Algorithms for Name Tagging**. In Proc. Human Language Technology conference - North American chapter of the Association for Computational Linguistics annual meeting (HLT-NAACL) Workshop on Computationally Hard Problems and Joint Inference in Speech and Language Processing, 2006.

106. Cynthia Rudin and Robert E. Schapire. **Margin-Based Ranking and an Equivalence Between AdaBoost and RankBoost**. Journal of Machine Learning Research, volume 10, pages 2193–2232, 2009.

  - Preliminary version: Cynthia Rudin, Corinna Cortes, Mehryar Mohri, and Robert E. Schapire. **Margin Based Ranking Meets Boosting in the Middle**. Proceedings of the Eighteenth Annual Conference on Learning Theory (COLT), pages 63 - 78, 2005.

2008 and before

107. Cynthia Rudin, Robert E. Schapire and Ingrid Daubechies. **Analysis of Boosting Algorithms Using the Smooth Margin Function**. Annals of Statistics, volume 35, number 6, pages 2723-2768, 2007.

  Preliminary material:

  - Cynthia Rudin, Robert E. Schapire, and Ingrid Daubechies. (2007) **Precise Statements of Convergence for AdaBoost and arc-gv**. In Proc. AMS-IMS-SIAM Joint Summer Research Conference: Machine Learning, Statistics, and Discovery, pages 131-145, 2007.
  - Cynthia Rudin, Robert E. Schapire, and Ingrid Daubechies. **Boosting Based on a Smooth Margin**. Proceedings of the Seventeenth Annual Conference on Computational Learning Theory, (COLT), pages 502-517, 2004.
  - Cynthia Rudin, Ingrid Daubechies, and Robert E. Schapire. **On the Dynamics of Boosting**. Advances in Neural Information Processing Systems (NIPS) 16, 2003.

108. Cynthia Rudin, Ingrid Daubechies, and Robert E. Schapire. **The Dynamics of AdaBoost: Cyclic Behavior and Convergence of Margins**. Journal of Machine Learning Research, 5 (Dec): 1557–1595, 2004.

  Preliminary material for this work appears partly within the NIPS paper below, and the open problem in COLT is from the JMLR paper:

- Cynthia Rudin, Ingrid Daubechies, and Robert E. Schapire. **On the Dynamics of Boosting**. Advances in Neural Information Processing Systems (NIPS) 16, 2003.
- Cynthia Rudin, Robert E. Schapire and Ingrid Daubechies. **Does AdaBoost Always Cycle?** JMLR: Workshop and Conference Proceedings, Published as a COLT Open problem, 2012.

109. Ryan Roth, Owen Rambow, Nizar Habash, Mona Diab, and Cynthia Rudin. **Arabic Morphological Tagging, Diacritization, and Lemmatization Using Lexeme Models and Feature Ranking**, The 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL/HLT), 2008.
110. Cynthia Rudin and Brian Spencer. **Equilibrium Ridge Arrays in Strained Solid Films**. Journal of Applied Physics, vol 86, pp 5530-5536, 1999.

Non-Peer-Reviewed Publications

111. Cynthia Rudin. **'The Marriage Pact' and the risks we take with data**, The News&Observer, February 28, 2021
112. Ezra Miller, Cynthia Rudin and Ingrid Daubechies. **Response to Letter to AMS Notices: Boycott collaboration with police**, September issue, AMS Notices, 2020. *We urge the mathematical community not to boycott interactions with police, but instead to work together to improve society.*
113. Matthew J Salganik, Lauren Maffeo and Cynthia Rudin. **Prediction, Machine Learning, and Individual Lives: an Interview with Matthew Salganik**. Harvard Data Science Review, 2020
114. Sarah Desmarais, Brandon Garrett, and Cynthia Rudin. **Risk Assessment Tools Are Not A Failed 'Minority Report'**. Perspectives, Law360, July 19, 2019, 5:50 PM EDT
115. Cynthia Rudin and David Carlson. **The Secrets of Machine Learning: Ten Things You Wish You Had Known Earlier to be More Effective at Data Analysis**. INFORMS TutORial, 2019.
116. Cynthia Rudin (with credit to Robin Smith). **Algorithms and Justice: Scrapping the Black Box.** The Crime Report, 2018.
117. **Analytic Research Foundations for the Next-Generation Electric Grid.** The National Academies Press, 2016. John Guckenheimer, Thomas Overbye (Co-chairs), and committee: Daniel Bienstock, Anjan Bose, Terry Boston, Jeffery Dagle, Marija D. Ilic, Christopher K. Jones, Frank P. Kelly, Yannis G. Kevrekidis, Ralph D. Masiello, Juan C. Meza, Cynthia Rudin, Robert J. Thomas, and Margaret H. Wright.
118. Cynthia Rudin, David Dunson, Rafael Irizarry, Hongkai Ji, Eric Laber, Jeffrey Leek, Tyler McCormick, Sherri Rose, Chad Schafer, Mark van der Laan, Larry Wasserman, Lingzhou Xue. **Discovery with Data: Leveraging Statistics with Computer Science to Transform Science and Society**. American Statistical Association whitepaper, `http://www.amstat.org/po cy/pdfs/B gDataStat st csJune2014.pdf`, 2014.
119. Cynthia Rudin and Kiri Wagstaff. **Machine Learning for Science and Society**, Machine Learning, (Introduction to the Special Issue on Machine Learning for Science and Society), volume 95, issue 1, April 2014, pp 1-9.
120. Cynthia Rudin. **Teaching "Prediction: Machine Learning and Statistics"**, Proceedings of the ICML Workshop on Teaching ML, 2012.
121. Peter Qian, Yilu Zhou, and Cynthia Rudin, **Proceedings of the 6th INFORMS Workshop on Data Mining and Health Informatics (DM-HI)**, editors, 2011.
122. Cynthia Rudin and Miroslav Dudĭk, **Lecture Notes for the AMS Short Course on Statistical Learning**, editors, includes contributions by Robert E Schapire, Lawrence Saul, Lisa Hellerstein, Adam Tauman-Kalai, and John Lafferty, 2007.

Grants

FAI: An Interpretable AI Framework for Care of Critically Ill Patients Involving Matching and Decision Trees, pending, NSF & Amazon, 2022. PI, with co-PIs Alex Volfovsky and Sudeepa Roy. 07/01/2022-06/30/2025. Amount of Award: $1,000,000

Center for Virtual Imaging Trials. NIH 1P41-EB028744-01A1, Co-lead for TRD3 PI-Ehsan Samei.

MedX: Interpretable Deep Learning Models for Better Clinician-AI Communication in Clinical Mammography, 11/1/2021-10/31/2022. Internal Duke Award, with Joseph Lo, Amount of Award: $50,000 total.

A Machine Learning Framework for Understanding Impacts on the HIV Latent Reservoir Size, Including Drugs of Abuse, National Institute on Drug Abuse (NIDA), PI, with co-I's David Murdoch, Nilu Goonetilleke and Nancie Archin. 9/30/2021-6/30/2026. Amount of Award: $2,284,375 total.

Research on Explainable AI on Multiple Models, Fujitsu, PI, 07/16/2021-7/15/2022. Amount of Award: $90,769.00 total

EAGER: Creating an Unsupervised Interpretable Representation of the World Through Concept Disentanglement, NSF, PI, 11/01/2021-10/31/2023. Amount of Award: $169,345 total

NSF Workshop on Seamless/Seamful Human-Technology Interaction, NSF, PI, 09/01/2021-08/31/2022. Amount of Award: $49,777 total

FAIR Data and Interpretable AI Framework for Architectured Metamaterials, DOE, co-PI, with PI Cate Brinson, and co-PI Chiara Daraio, 9/1/2020-8/31/2023. Amount of Award: $870,580 total

NRT-HDR Harnessing AI for Understanding & Designing Materials (aiM), FAIN 2022040, co-PI, with PI Cate Brinson, co-PIs David Banks, Stafano Curtarolo, Johann Guilleminot, 09/01/2020-08/31/2025. Amount of Award: $2,252,971 total.

NSF-HDR TRIPODS: Innovations in Data Science, CCF-1934964 co-PI, with PI Sayan Mukherjee, co-PI's Robert Calderbank, Rong Ge, and Jianfeng Liu, 10/1/2019-9/30/22. Amount of Award: $1.5 million total.

Utilizing Key Past Experiences from Large Datasets to Make Better Prediction in Multi-Class Settings, co-PI with Ramin Moghaddass, Amazon AWS Machine Learning Research Awards program, 2018. Amount of Award: $20,000.

An Integrated Nonparametric Bayesian and Deep Neural Network Framework for Biologically-Inspired Lifelong Learning (co-PI, with PI Katherine Heller, and other co-PI's Lawrence Carin, David Dunson, Nicolas Brunel, Tamara Broderick, Joshua Tenenbaum, Michael Jordan, Thomas Griffiths), DARPA, 2017-2020. Amount of Award: $4,388,119 total.

Collaborative Research Framework: Data: HDR: Nanocomposites to Metamaterials: A Knowledge Graph Framework, OAC-1835782 (co-PI, with PI Cate Brinson, and other co-PI's Chiara Daraio, Linda Schadler, Deborah McGuinness, and Wei Chen), NSF, 2018-2023. Amount of Award: $2,590,810 total.

Lord Foundation "Duke's Super Superresolution Team (Continued)" (Single PI) 2019-2020. Amount of Award: $24,000

Lord Foundation "Duke's Super Superresolution Team" (Single PI) 2018-2019. Amount of Award: $32,000

Duke Energy Initiative "Enabling Better Energy Decisions Through Better Interpretable Causal Inference Methods for Personalized Treatment Effects" Duke University Energy Initiative Energy Research Seed Fund (ERSF). (PI, with co-PI's Alex Volfovsky and Sudeepa Roy), 2018-2021. Amount of Award: $40,500 Stage I, $40,500 Stage II

QuBBD: Collaborative Research: Matching Methods for Causal Inference: Big Data and Networks, 1R01EB025021-01, sponsored by DHHS, PHS, NIH, NIBI&B" (co-Investigator, with PI Alexander Volfovsky, co-Investigators Sudeepa Roy and Allison Aiello). 2017-2020. Amount of Award: $849K ($514K to Duke)

Alfred P. Sloan Foundation "Interdisciplinary Energy Data Analytics Ph.D. Fellows Program," (co-PI, with Brian Murray, William Pizer, and Kyle Bradbury). 2018-2020. Amount of Award: $225,000

Laura and John Arnold Foundation "Interpretable Machine Learning for Pre-Trial Risk Analysis," 2017-2020. (Single PI) Amount of Award: $87,507

Duke Institute for Health Innovation "Pallialytics: Using analytics to inform a palliative care population health management intervention" (co-PI with Jonathan Fischer, Leslie Alabi, Eugenie Komives from Duke Medical). 2017-2018. Amount of award: $55,500

2016 Adobe Digital Marketing Research Award, "Compute-intensive causal machine learning models for finding customer segments" (Single PI). 2016. Amount of Award: $50K

MIT-Lincoln Labs "Adaptable, Interpretable, Machine Learning" (co-PI with Jonathan Su from MIT-LL), 2016-2019. Amount of Subaward: $542K

Xerox Research "Causal Inference-related Research Threads in the Prediction Analysis Lab, continued" (co-PI with Theja Tulabandhula), 2016. Amount of Award: $30K

DARPA "Foundations of Sequential Learning" (co-PI with Ron Parr and Kamesh Munagala), April 2016 - Jan 2017. Amount of Award: $242K

Philips "Algorithms for Interpretable Risk-Scoring" (Single PI), July 2017-December 2017. Amount of Award: $122K

Xerox Research "Causal Inference-related Research Threads in the Prediction Analysis Lab" (co-PI with Theja Tulabandhula), 2015. Amount of Award: $30K

Philips "Self-Learning Systems and User-Behavior Modeling" (Single PI), June 2015-May 2016. Amount of Award: $240K

Army Research Office "Uncertainty Quantification for Unobserved Variables in Dynamical Systems and Optimal Experimental Design" (Single PI), Spring 2015. Amount of Award: $50K

Big Data Seed Grant, MIT Big Data Initiative "Interpretable, Scalable and Causal Models from Machine Learning" (Single PI), Spring 2015. Amount of award: 1 semester RA

Accenture and MIT Alliance in Business Analytics "Creating a Model of the Usual State of a Machine" (Single PI), March 20 2015 - March 20 2016. Amount of award: $100,000

Wistron Corporation "Interpretable Predictive Models from Machine Learning" (Single PI), September 4, 2014 - August 31, 2016.
Amount of award $ 300,000

Accenture and MIT Alliance in Business Analytics "Big Data Analysis for Plant and Commercial Optimization," June 15 2013 - June 14 2014. Amount of award: $ 100,000

Ford Racing "Predictive Analytics for Racing" (Single PI),
September 1 2012 - August 31 2013. Amount of award: $75,000
September 1 2013 - December 31 2014. Amount of award: $75,000
December 31 2014 - December 31 2015. Amount of award: $157,000

Siemens Energy, CKI University Innovation Initiative. "CKI Proposal: Augmented Data-Driven Diagnosis using Physical Models" (Single MIT PI with co-PIs at Siemens), June 1, 2013 - May 31, 2016
Amount of award: $350,000

Siemens Energy, CKI University Innovation Initiative. "CKI Proposal: Incorporating Prediction Analysis into XHQ" (Single MIT PI with co-PIs at Siemens), June 1 2013 - May 31 2015.
Amount of award: 150,000 euro ≈ $200,000

Ford - MIT Alliance. "Develop advanced in-vehicle SYNC advertisement features to target drivers based on their context information, system interaction and past choices" (Single PI), July 1 2012 - August 31 2014.
Amount of award: $298,000

MIT Sloan Research Fund. "Predictive Models for Highly Imbalanced data," (Single PI), Amount of award: $20,000

Solomon Buchsbaum Research Fund. "A New Foundation for Statistical Decision-Making," June 2 2011-present. (Single PI), Amount of award: $50,000

NSF-CAREER IIS-1053407. "New Approaches for Ranking in Machine Learning," September 1 2011 – August 31 2018. Amount of award: $480,000

MIT Lincoln Labs. "A Knowledge Discovery Framework for Threat Identification (Phase III)" (co-PI with Shirish Ranjit from MIT-LL) 7/1/2013-6/30/2014.
Amount of award: $100,000

MIT Lincoln Labs. "A Knowledge Discovery Framework for Threat Identification (Phase II)" (co-PI with Shirish Ranjit from MIT-LL) 7/1/2012-6/30/2013.
Amount of award: $100,000

MIT Lincoln Labs. "A Knowledge Discovery Framework for Threat Identification" (co-PI with Shirish Ranjit from MIT-LL and Regina Barzilay from MIT), 6/1/2011- 5/31/2012.
Amount of award: $83,000

MIT Intelligence Initiative. "Combining human and machine predictions using "boosting" algorithms" (co-PI with Thomas Malone and Patrick Winston) November 1 2010 - October 1 2011.
Amount of award: 0.5 postdoctoral research fellowship

MIT Energy Initiative (MITEI) Seed Fund Program. "A Novel Framework for Electrical Grid Maintenance" (Single PI) May 17 2010 - May 16 2012.
Amount of award: $150,000

Ford - MIT Alliance. "Achieving Top Quality Ratings with Minimal Cost" (Single PI) July 1 2010 - June 20 2011.
Amount of award: $106,013

Con Edison Company of New York. "Manhole Events and Secondary System - Machine Learning Project"

Secondary System Project, Manhattan Backbone, December 1, 2011-December 31, 2013, Amount of subaward to MIT (Single PI): $300,129

Secondary System Project, Manhattan Corollary, May 1, 2011 - May 31, 2012, Amount of subaward to MIT (Single PI): $93,056

Update Manhattan Consolidated Database and Ranking Model through 2009, and Analyze 2004-2005 Secondary Inspections Data, July 1 2010 - December 31 2010, Amount of subaward to MIT (Single PI): $81,661

Phase 2, Application to B, Q, and X, December 1 2009 - May 31 2010, Amount of subaward to MIT (Single PI): $81,899

Phase 1, Application to B, Q, and X, January 2009 - June 2009, Co-PI , Amount: $464,127

Phase 4, July, 2008 - December, 2008, Co-PI , Amount: $413,947

Phase 3, January, 2008 - June, 2008, Co-PI , Amount: $347,104

Phase 2, August, 2007 - December, 2007, Co-PI, Amount: $486,296

Phase 1, March, 2007 - July, 2007, Co-PI , Amount: $339,251

National Science Foundation. Postdoctoral Research Fellowship in Biological Informatics, Grant DBI-0434636, March 2005 - February 2007.
Amount of award: $120,000

Honors (not including awards listed with papers above)

2022 Squirrel AI Award for Artificial Intelligence for the Benefit of Humanity from the Association for the Advancement of Artificial Intelligence (AAAI). This is the most prestigious award in artificial intelligence. This award, similar

only to world-renowned recognitions, such as the Nobel Prize and the Turing Award, carries a monetary reward at the million-dollar level.

Guggenheim Fellow, 2022

Fellow of the Association for the Advancement of Artificial Intelligence (AAAI), 2022-present

Thomas Langford Lecture Award, Duke University, 2019-2020

Fellow of the American Statistical Association, 2019-present

Fellow of the Institute of Mathematical Statistics, 2019-present

Duke AI for Art competition (with students Alina Barnett, James Hoctor, Chaofan Chen, and Oscar Li), second place, 2019

Faculty associate, Berkman Klein Center for Internet and Society at Harvard University, 2015-2019.

2016 Adobe Digital Marketing Research Award, 2016

Named by Businessinsider.com as one of the 12 most impressive professors at MIT in 2015

National Science Foundation CAREER Award, 2011

"Top 40 Under 40" business school professors of *Poets & Quants*, 2015. (Published in Forbes magazine.)

Nominated for Outstanding UROP Mentor Award, UROP (undergraduate research opportunities) Program, MIT, 2012

Nominated for 2012 Sloan Excellence in Teaching Awards, MIT, 2012

Second Place for Phase 1 in the ICML Exploration and Exploitation 3 Challenge, 2012. Goal is to design a recommender system with high click through rates for Yahoo! Front Page News Article Recommendations teammates: Virot Ta Chiraphadhanakul and Edward Su

National Science Foundation Postdoctoral Research Fellowship in Biological Informatics, 2005-2007

University at Buffalo College of Arts and Sciences Outstanding Senior Award in Sciences and Mathematics 1999, one per year at the university (also Department of Physics Outstanding Senior 1998, Department of Mathematics Outstanding Senior 1999, Department of Music Outstanding Senior 1999)

Barry J. Goldwater Scholarship, 1997-1998

State University of New York Chancellors Award for Student Excellence 1997, 1999

Dr. Stanley T. Sekula Memorial Scholarship, University at Buffalo Physics Department, 1996, 1997

Hildegarde F. Shinners Prize, 1999, Phi Beta Kappa award to recognize mathematics thesis and music thesis

Phi Beta Kappa, inducted 1997

UB Music Department Scholarships, 1994, 1995, 1996

Intellectual Breath and Liberal Knowledge Award, UB Honors Program, 1999

Second Place-National Winner of the 1993 Young Inventors' and Creators' Competition in the Copyright Category of Popular Music Composition, sponsored by the Foundation for a Creative America. (This came with a congratulatory letter signed by Vice President Al Gore!)

Media (Selected)

Squirrel AI Award:

"Duke Computer Scientist Wins $1 Million Artificial Intelligence Prize, A 'New Nobel' " AAAI press release, October 12, 2021

"Duke Professor Recognized for Clarifying AI Decision Making" by John McCormick, Wall Street Journal, Lead article in AI section, October 14, 2021

"She won a $1 million prize for predicting which manholes would explode" by James Barron, New York Times, New York Today section, October 12, 2021 (This title is somewhat misleading.)

Random Interpretability Interviews:

What Do Conspiracy Theories And AI Explainability Have In Common? Kareem Saleh, Forbes, May 4, 2021

"Meet Cynthia Rudin–A Champion of Interpretable Machine Learning," Sam Behseta & Michelle Dunn, Chance Magazine, American Statistical Association, April 22, 2020

"Rise of Robot Radiologists," by Sara Reardon, Nature, Innovations In, December 18, 2019

BBC radio, Digital Planet, December 10, 2019 (discussing "This Looks Like That")

"Machine vision that sees things more the way we do is easier for us to understand" MIT Technology review, Artificial Intelligence, Dec 6 2019

Computer Vision:

"Accurate Neural Network Computer Vision Without the 'Black Box': Duke team disentangles neural networks to understand how they see the world," December 15. 2020

"Artificial Intelligence Makes Blurry Faces Look More Than 60 Times Sharper: This AI turns even the blurriest photo into realistic computer-generated faces in HD": by Robin A. Smith, Duke Today, June 12, 2020 - published on DukeToday, covered by Newsweek.com, Independent.co.uk and other media outlets.

"Making Blurry Faces Photorealistic Goes Only So Far," by Mark Anderson. IEEE Spectrum, June 23 2020

Crime Data Mining:

"Possible to predict recidivism? Here's how...," The Docket, MSNBC (Live TV), Tuesday May 19, 2015.

"Crime-Fighting Computer Code from Cambridge Police and MIT," *WBUR Radio Boston* (National Public Radio), Tuesday August 13 2013

"Cambridge police look at math to solve crimes," *Boston Globe*, Metro Section, front page, Sunday August 4 2013

"Statistician enlisted to fight crime by numbers," *The Times of London*, US & Americas section, Tuesday August 6 2013

"Predictive Policing: Using Machine Learning to Detect Patterns of Crime," *WIRED Innovation Insights*, August 22, 2013

Meetings Analysis:

"At Work: Just Say 'Yeah'," *Wall Street Journal*, Business section, on page B8 in the U.S. edition, June 19, 2013

"How to be effective at meetings? Say 'yeah'," *Toronto Star*, Business section, June 28, 2013

"Researchers discover the key to persuasion," *ABC News*, consumer report blog / business, June 24, 2013

Energy Grid:

My work discussed in book The Alignment Problem: Machine Learning and Human Values 1st Edition by Brian Christian, W. W. Norton and Company, 2020.

"Why Manhole Explosions Happen in the Summer," *NBC News*, Business/Energy section, August 19, 2015

"New York's Exploding Manhole Covers Pose Unexpected Winter Hazard," *Reuters*, appeared in *New York Times*, and *MSN.com*, February 28, 2015

My work discussed in book Big Data: A Revolution that Will Transform how we Live, Work, and Think, by Victor Mayer-Schönberger and Kenneth Cukier, Houghton Mifflin Harcourt Publishing Company, 2013

Analytics Magazine, INFORMS. Headlines: Innovative Applications in Analytics Award, April 18, 2013

"Machine Versus Manhole," *ScienceNews, U.S. News and World Report, WIRED Science, Slashdot, Discovery News / Discovery Channel*, and others, July 8-9 2010

Information Retrieval / Building New Search Engines:

Radio segment about my work on Growing a List. "A New Way to Google," *Boston Public Radio*, show on innovation at 12:45pm-1pm, hosted by Kara Miller, October 9, 2012

Health and Interpretable Predictive Models:

"No more excuses. Make data more accessible." Washington Post, Opinions, part of collection "We need smart solutions to mitigate the coronavirus's impact. Here are 46." June 18, 2020

"How Can Doctors Be Sure A Self-Taught Computer Is Making The Right Diagnosis?," by Richard Harris, All Things Considered, NPR, April 1, 2019

"Do You Zone Out? Procrastinate? Might Be Adult ADHD," by Rebecca Hersher, NPR, April 5 2017

"Algorithms Learn From Us, and We've Been Bad Parents," by Bahar Gholipour, Mach Technology, NBC News, March 10 2017, 2:17 PM ET

"New Computer Tool Can Predict Dementia From Your Simple Drawings" Popular Science, August 13, 2015

"Digital Pen is Better Dementia-Prediction Tool than a Doctor" WIRED Magazine, August 17, 2015

"Computers that teach by example: New system enables pattern-recognition systems to convey what they learn to humans." MIT News (also front page of MIT main website), December 5-10, 2014

"New Statistical Model Lets Patient's Past Forecast Future Ailments," Science News section, *Science Daily*, June 9, 2012

Other topics:

Scientific Sonnets: Duke team wins competition for poetry-generating algorithm, Duke Chronicle, Dec 27, 2018

Article discussing the whitepaper effort I led: AmstatNews: The Membership Magazine of the American Statistical Association News. Cover Story, *Statistical Scientists Advance Federal Research Initiatives*, July 1, 2014.

Article about my work: "How to Improve Product Rankings," *Businessweek*, B-School Research Briefs section, March 9, 2012

Professional Societies and Government Committees

Executive Committee Member, ACM SIGKDD, 2021-present

Associate Director, Statistical and Applied Mathematical Sciences Institute (SAMSI), 2018-2021

Chair of Committee to Choose the First Editor-in-Chief of the new INFORMS Journal on Data Science, INFORMS, 2019-2020

Judge for Edelman Award, INFORMS 2019-2020, 2020-2021

Technology Strategy Committee, INFORMS, 2019-present

Member of Committee for Computing Community Consortium (CCC) for Interaction for AI / Roadmap for Future AI, January 2019

Chair, Section on Statistical Learning and Data Mining, American Statistical Association, 2017-2018.

Member of Committee on Applied and Theoretical Statistics (CATS), National Academies of Sciences, Engineering, and Medicine, 2016-present

Member of Committee on Law and Justice (CLAJ), National Academies of Sciences, Engineering, and Medicine, 2017-present

Councilor of the AAAI. 2017-2020

Chair, INFORMS Data Mining Section, 2015-2016 (Vice Chair for 2014-2015, Council Member 2017-2018, Council Member, 2011- 2013).

Member of Committee on Analytical Research Foundations for the Next-Generation Electric Grid, and author of consensus report: "Analytical Foundations for the Next Generation Electric Grid", National Academy of Sciences, Engineering and Medicine, 2014-2016.

Bureau of Justice Assistance Criminal Justice Technology Forecasting Group (BJA CJTFG), United States Department of Justice, 2014-2016.

DARPA Information Science and Technology (ISAT) study group (faculty advisory board of DARPA), 2014-2018.

American Statistical Association Committee on Funded Research, 2015-2018.

MIT Energy Initiative, Energy Education Task Force, 2013-2014.

Activities

Events Organized

FAIF: Fair AI in Finance, NeurIPS workshop, co-organizer, 2020.

Self-Supervised Learning – Theory and Practice, NeurIPS workshop, co-organizer, 2020.

Triangle Machine Learning Day, lead organizer, co-organized with David Banks, Jeremy Freeman, and Ted Enamorado, September 20, 2019.

SAMSI Deep Learning Program, co-organizer and leader of working group, 2019

Conference co-Chair, Conference on Statistical Learning and Data Science / Nonparametric Statistics, co-Chair with Annie Qu, American Statistical Association, June 4-6, NYC, 2018

Triangle Machine Learning Day, lead organizer, co-organized with Jade Vinson, Richard Lucic, and Kirsten Shaw, April 3, 2018.

Member of planning committee, SAMSI Program on Statistical, Mathematical, and Computational Methods for Precision Medicine (2016-2018)

TAMALE: Toolkit of Algorithms for Machine Learning, DARPA ISAT workshop, co-organized with Margo Seltzer, March 2018.

Judge for INFORMS Data Mining Best Paper competitions, INFORMS, 2018.

Member of organizing committee, Conference on Statistical Learning and Data Science, UNC Chapel Hill. 2016.

What if: Machine Learning Models for Causal Inference, DARPA ISAT workshop, co-organized with Dustin Tingley, February 2016.

The Cassandra Problem: Building Trust in Predictive Models, DARPA ISAT workshop, co-organized with Carla Brodley and Stephen Boyd, April 2015.

Invited Session: The Fifth V in "Big Data" is *Variables*, co-organizer with Tyler McCormick, Joint Statistical Meetings, 2015.

Topic Contributed Session: Predictive Policing, organizer, Joint Statistical Meetings, 2014.

Judge for Statistical Learning and Data Mining Best Student Paper competition, American Statistical Association, 2015.

Judge for 2015 INFORMS Innovative Applications in Analytics Award, 2014-2015.

Workshop on Data Analytics: Challenges in Big Data for Data Mining, Machine Learning and Statistics organizer, MIT CSAIL Big Data, March 26, 2014.

The ISBIS (International Society of Business and Industrial Statistics) 2014 and SLDM (Statistical Learning and Data Mining section of the American Statistical Association) Meeting on Data Mining in Business and Industry, Program Committee, June 9-11, 2014.

Statistical Analysis and Data Mining (a journal of the American Statistical Association), committee to choose the next editor-in-chief, 2014-2015.

ECML/PKDD 2013 Workshop on "DARE: Data Analytics for Renewable Energy Integration", Technical program committee member, 2013.

Workshop on Hospital Readmission Prediction and Clinical Risk Management (HRPCRM) at IEEE International Conference on Healthcare Informatics (ICHI) 2013, program committee member, organizers are John Cromwell and Si-Chi Chin

Session on Smart Grid Data Analytics (SGDA) at International Conference on Smart Grid and Clean Energy Technologies (ICSGCE) 2013, co-chair with Zeyar Aung

Dagstuhl Seminar on Preference Learning, co-organizer, 5 day seminar, 45 participants, Germany, March 3th to March 7th, 2014.

IMS/ASA Spring Research Conference, organizer of Machine Learning session, Harvard University, June 14th, 2012.

New England Machine Learning Day, Co-organizer, Microsoft Research New England, May 16th, 2012.

New England Statistics Symposium (NESS), organizer of Machine Learning session, Boston University, April 21, 2012.

Collective Intelligence 2012 (CI 2012), Local Arrangements Chair, 2012

INFORMS Data Mining and Health Informatics Workshop (DH-MI), co-organizer, 2011

MIT Energy Initiative, Organizing Committee for the MITEI Seminar Series, member, Fall 2010- Spring 2012

AMS Short Course on Aspects of Statistical Learning, organizer, 2007 AMS joint math meetings, New Orleans, January 3-4, 2007

AMS/AWM/MAA Special Session on Mathematical Results and Challenges in Learning Theory, session organizer, AMS joint math meetings, San Antonio Texas, January 12-15, 2006

Program for Women in Mathematics, Institute for Advanced Study, Program Committee Member, 2003-2006. Women in Science Seminar organizer, 2004, 2005, 2006. TA for the Wavelets course in 2002. Discussion group organizer, 2001. Research seminar speaker 2003. Poster session organizer 2003, Panel discussion organizer 2003, Panelist 2007

PACM Conference Princeton University, organizer, 2002-2004, speaker 2005

Editorial Responsibilities

Associate Editor for Journal of Quantitative Criminology, 2021 - present

Associate Editor for Harvard Data Science Review, 2019 - present

Associate Editor for Management Science, in the Big Data Analytics department, 2018 - present

Associate Editor for INFORMS Journal on Data Science, 2020 - present

Action editor: 2013 - 2017, (Editorial board member: June 30 2010 - June 30 2013), Machine Learning Journal

Action Editor: Statistical Analysis and Data Mining (SAM, a journal of the American Statistical Association), 2012-2017

Editor for Special Issue on Sports Analytics for Statistical Analysis and Data Mining (SAM, a journal of the American Statistical Association), co-editor with Theja Tulabandhula, 2015

Editorial Board, Journal of Artificial Intelligence Research (JAIR), July 2014 - June 2017

Member of Guest Editorial Board for ECMLPKDD 2014 journal track

Editor for Special Issue on Machine Learning for Science and Society, co-editor with Kiri Wagstaff, 2012-2013

Editorial board member: Journal of Machine Learning Research, 2012 - present

Program Committee Memberships

NIPS Workshop on Machine Learning for Healthcare ML4HC (2016, 2017, 2018), ECML-PKDD Workshop on Social Good (2016), ICML Workshop on Human Interpretability for Machine Learning (2016), ECML-PKDD Workshop on Data Science for Social Good (2016), SDM senior pc (2016), ICDM area chair (2015), AAAI senior pc (2016), NIPS area chair (2015), Visual Data Science (2015), IJCAI senior pc (2015), ICML area chair (2015) ICDM area chair (2014), IEEE Big Data (2013), NIPS area chair (2013), AAAI (2013), ICML area chair (2013), NIPS area chair (2012), ECML-PKDD (2012), ICML (2012), COLT (2011), IJCAI (2011), ECML-PKDD (2010), ICML (2009), ICML (2006), AAAI (2006), KDD (2005)

Referee Assignments

Journal Reviewing:

Biometrics, Harvard Data Science Review, Journal of Artificial Intelligence Research (JAIR), Journal of Quantitative Criminology (JOQC), Data Mining and Knowledge Discovery (DAMI), Journal of Machine Learning Research (JMLR), Machine Learning Journal, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Mathematical Programming, Journal of Computational Statistics and Data Analysis, Communications on Pure and Applied Mathematics (CPAM), Management Science, International Journal of Renewable Energy Research (IJRER), Annals of Operations Research, Recent Patents on Computer Science, The Lancet, Scientific Reports - Nature, Science Advances, Decision Support Systems

Conference and Other Reviewing: NeurIPS/NIPS (Neural Information Processing Systems) 2020, 2019, 2018, 2017, 2011, 2010, 2009, 2008, 2007. IEEE Computer 2019, ICDM (International Conference on Data Mining) 2011, Grant Proposal Review for Natural Sciences and Engineering Research Council of Canada 2016, Book Proposal for Springer 2013, COLT (Conference on Learning Theory) 2012, COLT 2010, COLT 2005, AISTATS (Conference on Artificial Intelligence and Statistics) 2012, AISTATS 2010, ICMLA (International Conference on Machine Learning and Applications) 2011, Book Proposal for Manning Publications 2010, ALT (Algorithmic Learning Theory) 2010, ICML (International Conference on Machine Learning) 2010, NIPS Ranking Workshop 2009, ACM SIGKDD (Conference on Knowledge Discovery and Data Mining) 2009, Applied and Computational Harmonic Analysis Journal, 2006, Conference on Machine and Statistical Learning: Prediction and Discovery 2006

Presentations

Keynotes/Invited/Plenary for Notable Conferences

SDM (SIAM International Conference on Data Mining), Keynote, April 29, 2022

AAAI, Plenary talk for Squirrel AI Award, February 24, 2022

Nobel Conference, Gustavus Adolphus College, October 4-6, 2021

Machine Learning in Science & Engineering (MLSE), Columbia University, December 14, 2020

Data Science, Statistics & Visualization (DSSV), SAMSI, July 29, 2020

Conference on Digital Experimentation (CODE), MIT, November 2, 2019

IEEE International Conference on Data Science and Advanced Analytics (DSAA), Washington DC, October 7th, 2019

25th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), Anchorage, Alaska, August 8th, 2019

Advances in Decision Analysis, INFORMS Decision Analysis Society, (keynote), Bocconi University, Milan, June 19, 2019

Women in Data Science Cambridge Conference (organized by Harvard, MIT, Microsoft Research New England, Stanford) March 4, 2019

The European Conference on Machine Learning & Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD), Dublin, September 11, 2018

Females Excelling More in Math, Engineering, and Science (FEMMES), keynote for capstone event, 200 girls grades 4-6, Duke University, February 17, 2018

International Symposium of Artificial Intelligence and Mathematics (ISAIM), January 3, 2018

Machine Learning for Healthcare (MLHC), August 18, 2017

Artificial Intelligence and Statistics (AISTATS), April 22, 2017

The Frontiers of Machine Learning, Sackler Forum on Machine Learning, National Academy of Sciences and The Royal Society, Jan 31-Feb 1, 2017

Fairness, Accountability and Transparency in Machine Learning (FAT-ML), November 18, 2016

Predictive Applications and APIs (PAPIs), October 2016

Discovery Science, Banff, Canada. October 5, 2015.

20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), August 24-27, 2014.

Nobel Conference, Gustavus Adolphus College, October 5-6, 2021

Other Conference Invited Presentations

(This list does not include regular or invited talks at annual meetings such as JSM, INFORMS, SLDM/SLDS, American Society for Criminology, Atlantic Causal Inference Conference, which I participate in most years.)

Interpretability in Artificial Intelligence Workshop, Banff International Research Station, Keynote, May 3, 2022

Discover Experiential AI, The Institute for Experiential AI, Northeastern University, Panelist, April 6, 2022

Walmart AI Summit, Plenary speaker, April 6, 2022

INFORMS Optimization Society Conference, Plenary speaker, Greenville, SC, March 14, 2022

5th Annual Swedish Symposium on Deep Learning / 39th Annual Swedish Symposium on Image Analysis, March 14, 2022

Stu Hunter Research Conference on Statistics and Statistical Engineering, 2022, keynote, Duke University, March 8, 2022

AAAI Workshop on Interactive Machine Learning, keynote, February 28, 2022

AAAI Explainable Agency in Artificial Intelligence Workshop, Keynote, February 28, 2022

Government of Canada Data Conference 2022: Driving Data Value and Insights for All Canadians, Panelist, February 23, 2022

Fidelity AI Day, keynote, February 16, 2022

The Where and Why of Explainable AI, Science and Technology Expert Partnership Advances in Explainable AI Workshop, MITRE Labs, panelist, January 11, 2021

5th Joint International Conference on Data Science & Management of Data (CODS-COMAD 2022), keynote, India (virtual), January 7, 2022

HICSS-55: SWT on Future of Human Work: Harnessing the Power of Augmented Intelligence and Augmented Cognition, panelist, January 3, 2022

Interpretability, Safety, and Security in AI, The Alan Turing Institute, December 14, 2021

Human-Centered AI (HCAI) Workshop at NeurIPS, keynote, December 13, 2021

Machine Learning & Supply Chain Management Workshop, TRIPODS-X, December 13, 2021

Forward Summit, Puerto Rico Science, Technology & Research Trust, December 10, 2021

2021 IEEE Symposium Series on Computational Intelligence (SSCI), Deep Learning Track, keynote, December 5, 2021

Gillmore Symposium: Explainable, Interpretable AI: the Future of Investment Management, Warwick University, November 19, 2021

Conference on Non-traditional Data, Machine Learning and Natural Language Processing in Macroeconomics, Bank of Canada, panelist, November 18, 2021

DARPA DSO Futures meeting, plenary speaker, November 18, 2021

Duke in DC discussion on "The Equitable, the Ethical and the Technical: Artificial Intelligence's Role in the U.S. Criminal Justice System," November 15, 2021

McCombs School of Business' Center for Analytics and Transformative Technologies 2021 Global Analytics Conference, keynote, November 12, 2021

Artificial Intelligence in Consumer Finance: Defining and Insuring Fairness, Federal Reserve Bank of Philadelphia and Federal Reserve Bank of Cleveland, November 9, 2021

AAAI Fall Symposium, keynote, November 5, 2021

Advanced Analytics: New Methods and Applications for Macroeconomic Policy, organized by the Bank of England, the European Central Bank, King's College London and King's Business School, keynote, November 4, 2021

Methodological Approaches for Whole Person Research Workshop, NIH, September 29-30, 2021

CNRS School on Explainability, September 30, 2021

iMIMIC Workshop on Interpretability of Machine Intelligence in Medical Image Computing at MICCAI 2021, September 27 2021

Explainable AI Virtual Workshop, Caltech, September 23, 2021

AI Meets Regulators Symposia, AI for Health at Imperial College, September 21, 2021

ICCAI'21 International Conference on Complex Acute Illness, keynote, September 10, 2021

FICO Mastermind Event, Aug 30 – Sept 19, 2021

ARES / CD-MAKE 2021 Conference, keynote, August 20, 2021

ELLIS workshop on Causethical ML, Invited talk, July 26, 2021

ICML Workshop on Theoretic Foundation, Criticism, and Application Trend of Explainable AI, Invited talk, July 23, 2021

CVPR 2021 Tutorial on Interpretable Machine Learning

Forecasting the future for sustainable development, Centre for Excellence and Transdisciplinary Studies, hosted by OECD, keynote, June 17, 2021

Analytics Summit 2021, University of Cincinnati, June 8, 2021

Responsible Machine Learning, keynote, North Carolina State University, June 4, 2021

ICLR Workshop on Responsible AI, May 7, 2021

Bringing Artificial Intelligence to the Bedside, Purdue University Workshop, April 23, 2021

Trustworthy Automated Decision Making (ETAPS 2021 Workshop), keynote, March 28, 2021

Explainable AI Mini-Summit, Re-Work, February 17, 2021

Panel: Bias in AI: How scientists are trying to fix it?, Intuit, February 15, 2021

Florida Women in Mathematics Day, Florida Atlantic University, keynote, February 13, 2021

IEEE EMBS Forum on Data Science and Engineering in Medical Imaging, Symposium #1: Grand Challenges in Data Science and Engineering in Healthcare: Medical Imaging, February 10, 2021

The WACV 2021 Workshop on Explainable & Interpretable Artificial Intelligence for Biometrics (xAI4Biometrics Workshop 2021), keynote, January 5, 2021

NeurIPS Workshop on Broader Impact of AI, December 12, 2020

Toronto Machine Learning Summit, November 19, 2020

Advancing Analytics 2020, National Conference, Institute of Analytics Professionals of Australia, November 17, 2020

The Machine Learning Conference (MLconf), November 6, 2020

Explainable AI Planning (XAIP) @ ICAPS, October 21, 2020

Symposium on Artificial Intelligence for Learning Health Systems (SAIL), Presymposium, October 21, 2020.

Advances in Interpretable Machine Learning and Artificial Intelligence (AIMLAI), keynote, workshop at CIKM, October 20, 2020

Workshop on Credit Card Lending and Payments, Federal Reserve Bank of Philadelphia, September 17, 2020

National Health Symposium, Johns Hopkins Applied Physics Laboratory, September 14, 2020

Workshop on New Directions in Optimization, Statistics and Machine Learning, Institute for Advanced Study, April 16, 2020

TutORial, INFORMS, October 21, 2019

Debugging Machine Learning Models: ICLR 2019 workshop, May 6, 2019

Safe Machine Learning: Specification, Robustness and Assurance: ICLR 2019 Workshop, May 6th, 2019

FEAP-AI4Fin 2018 : NeurIPS 2018 Workshop on Challenges and Opportunities for AI in Financial Services: the Impact of Fairness, Explainability, Accuracy, and Privacy, invited talk, December 7, 2018

NeurIPS 2018 Workshop on Critiquing and Correcting Trends in Machine Learning, spotlight talk, December 7, 2018

Nicholas Institute, Duke University, invited commenter on deep learning session, October 5, 2018

3rd International Workshop on Biomedical Informatics with Optimization and Machine Learning (BOOM), invited talk, IJCAI workshop, June 13, 2018

Workshop on Human Interpretability in Machine Learning (WHI), invited talk, ICML workshop, June 17, 2018

CVPR NTIRE 2018 New Trends in Image Restoration and Enhancement workshop and challenges on super-resolution, dehazing, and spectral reconstruction, invited talk, representing our competition entry on superresolution, June 18, 2018

Triangle Machine Learning Day, invited talk, April 3, 2018

4th Annual Morgan Stanley Quantitative Equity Research Conference, invited talk, November 17, 2016

Japan-America Frontiers of Engineering Symposium, National Academy of Engineering, invited talk, Beckman Center, Irvine CA, June 16, 2016

Criminal Justice in the Age of Big Data, panelist. Harvard Kennedy School. November 13, 2015

Duke University Workshop on Sensing and Analysis of High-Dimensional Data, Duke University, July 27-29, 2015

Operations, Technology and Information Management Research Camp, Johnson School of Management, Cornell University. June 23, 2015

MIT Big Data Workshop – "From Data to Insights". May 21, 2015.

MIT Media Lab – DARPA ISAT Workshop on Computational Health. April 20, 2015

Workshop on Big Data and Statistical Machine Learning. (Part of the Thematic Program on Statistical Inference, Learning, and Models for Big Data.) Fields Institute, Toronto, January 29, 2015

Bringing Social Science Back In: The 'Big Data' Revolution and Urban Theory, Radcliffe Institute, Harvard University, December 15-16, 2014

Understanding and Improving Cities: Policy/Research Partnerships in the Digital Age, Invited Speaker/Panelist, District Hall, Boston, December 12, 2014

Weathering the Data Storm: The Promise and Challenges of Data Science, Third Annual Symposium on the Future of Computation in Science and Engineering, Harvard University, Jan 26, 2014.

New England Machine Learning Day, May 1, 2013

Exploration and Exploitation 3, Presentation for Second Place team in Data Mining Contest, ICML Workshop, June 2012

IMA Workshop on User-Centered Modeling, May 2012

IMA Workshop on Machine Learning, March 2012

ICML Workshop on Global Challenges, International Conference on Machine Learning, July 2, 2011

American Institute of Mathematics (AIM) workshop on the Mathematics of Ranking, August 16-20, 2010

DIMACS/CCICADA Workshop on Statistical Issues in Analyzing Information from Diverse Sources, Rutgers University, May 6-7, 2010

New England Statistics Symposium (NESS), Harvard University, April 17, 2010

International Utility Working Group: Workshop on Computer-Aided Lean Management (CALM), Columbia University, April 16, 2008

Conference on Machine and Statistical Learning: Prediction and Discovery, AMS-IMS-SIAM Summer Research Conferences in the Mathematical Sciences, invited by organizers Joe Verducci, Xiaotong Shen, and John Lafferty, Snowbird, Utah, June 25-29th, 2006

FOCM 2005 Foundations of Computational Mathematics, Workshop 4 on Learning Theory, Santander Spain. Invited by organizer Steve Smale, 2005

Machine Learning Summer School, Workshop on the Dynamics of Learning, TTI-C, Chicago, invited by organizer Steve Smale, May 16-26, 2005

Invited Presentations at Universities, Societies, and Research Labs

Fields Machine Learning Seminar, Fields Institute for Research in Mathematical Sciences University of Toronto, April 25, 2022
Data Valorization Institute (IVADO) in Montreal, Zooming in on Multidisciplinary AI, April 21, 2022
Michigan Institute for Data Science Seminar Series, April 14, 2022
International Conference on Foundations and Applications of AI, Peking University, April 8, 2022
UCSB Center for Responsible Machine Learning Distinguished Lecture Series, UC Santa Barbara, April 8, 2022
"I Can't Believe It's Not Better" Seminar Series, April 7, 2022
Elon University Analytics Day, March 24, 2022
Cambridge Centre for AI in Medicine Seminar, Cambridge UK, March 16, 2022
Janelia Research Campus Computation & Theory Seminar, Janelia Farms, March 10, 2022
The Alan Turing Institute, Human-Machine Teams Seminar, February 25, 2022
Institute for Experiential AI, Northeastern University, Bay Area, Distinguished Speaker Series February 23, 2022
University of Virginia, AI and Machine Learning Seminar, February 4, 2022
Institute for Assured Autonomy Speaker Series, Johns Hopkins University, February 3, 2022
London School of Economics & Political Science, Data Science Seminar, January 31, 2022
Roche Advanced Analytics Network, Roche/Genentech, January 26, 2022
Image Guided Cancer Therapy Seminar, MD Anderson Cancer Center, January 19, 2022
Perspectives in Mathematical Sciences, Dr. F.C. Kohli Centre of Excellence, Chennai Mathematical Institute, January 19, 2022
Trustworthy ML, Fireside Chat, December 16, 2021
The Myth of "Explainable" AI and why "Interpretable" AI is the Answer, Collective AI Podcast, December 15, 2021
AI for Good, Accelerating the United Nations Sustainable Development Goals, Seminar Series, International Telecommunication Union, December 9, 2021
Montreal Speaker Series in the Ethics of AI, December 9, 2021
Chalmers AI Talks, Chalmers University in Sweden, December 8, 2021
Teaching Youth to Build and Deploy Responsible AI for Justice Workshop, Technovation, December 7, 2021
ECE Distinguished Speaker Series, Rice University, December 6, 2021
Oxford Women in Computer Science Distinguished Speaker Series, November 9, 2021
Smarsh Advanced AI panel on Driving Adoption of AI, November 1, 2021
Responsible Modelling in Uncertain Times — CEST-UCL Seminar series, Keynote, November 3, 2021
RBCDSAI LatentView Colloquium, Robert Bosch Centre for Data Science and AI at IIT-Madras, October 29, 2021
TOM Seminar, Harvard Business School, October 23, 2021
Optimization and Data Science Community Seminar, Exxon, October 7, 2021
Neyman Seminar, Berkeley Statistics Department, September 15, 2021
Trustworthy ML Initiative Seminar, Harvard, July 8, 2021
Fidelity AI Seminar, June 29, 2021
University of Basel, Computational Seminar Series, Computational Biology core program of the Biozentrum, Basel, Switzerland, May 31, 2021
University of Oxford, Oxford Computational Statistics and Machine Learning (OxCSML) Seminar, May 14, 2021
FDIC Center for Financial Research Seminar, May 11, 2021
EPFL Statistics Seminar, May 7 2021
Distinguished Seminars on Explainable AI, connected to ERC project "XAI Science and technology for the eXplanation of AI decision making," April 20, 2021
Frontiers of Big Data, AI, and Analytics, Virtual Seminar Series (centered in Australia), April 14, 2021
Iowa State University, Theoretical and Applied Data Science Seminar, April 8, 2021
National Institute of Cancer, Biostatistics Branch, Division of Cancer Epidemiology and Genetics, Seminar, April 7, 2021
New Jersey Institute of Technology Data Science Seminar, March 10, 2021
Metron/George Mason University AI/ML Seminar Series, March 10, 2021
NeEDS Mathematical Optimization Seminar Series, Network of European Data Scientists, February 8, 2021
Brown University, Division of Applied Mathematics Colloquium, February 4, 2021
Texas A&M Institute of Data Science, TAMIDS Seminar, January 30, 2021
Bias^2 Seminar, Harvard Data Science, January 28, 2021
University of Oxford, Computer Science, OATML group, January 26, 2021
University of California, Santa Cruz, Winter Seminar Series, Department of Statistics, January 25, 2021
Verizon Media Research Day, keynote, December 17, 2020
University of Iowa, Department of Business Analytics Seminar Tippie College of Business, December 11, 2020

Duke Computing Roundtable, December 10, 2020
Energy Data Analytics Symposium: Transforming Energy Systems with Data Science Techniques, Duke University, December 9, 2020
Alberta Machine Intelligence Institute, University of Alberta, December 4, 2020
North Carolina Agricultural and Technical State University (NCAT), Seminar, organized by Student Leadership Council and Faculty of the ACIT Institute and TECHLAV Center, November 20, 2020
University of Pittsburgh, Statistics Department Seminar, October 27, 2020
University of California, Davis, Mathematics of Data and Decision in Davis Seminar, October 20, 2020
University of Toronto, MIE Distinguished Seminar Series, Department of Mechanical and Industrial Engineering, October 2, 2020
University of Colorado, Denver, Biostatistics Seminar, September 16, 2020
Online Causal Inference Seminar, August 25, 2020
Tutorial on Interpretable Machine Learning, Joint Statistical Meetings, August 3, 2020
TRIPODS Seminar Series, TRIPODS DATA-INSPIRE Institute, DIMACS & Rutgers CS/Math/Statistics, July 17, 2020
Melbourne Centre for Data Science Seminar Series, June 25, 2020
Decision Making in an Uncertain World, Seminar Series, INFORMS Stochastic Programming Society, June 12, 2020
Center for Human-Compatible AI (CHAI) Seminar, Berkeley, June 10, 2020
Boeing Research Seminar, June 5, 2020
Research Triangle Institute (TRI) Seminar, April 21, 2020
Carnegie Mellon Artificial Intelligence Seminar, April 7, 2020
GSS/SSS webinar program (offered by the American Statistical Association's Government Statistics Section and Social Science Section, February 13, 2020
Ingrid Daubechies Lecture in Computer Science, Duke University, January 21, 2020
North Carolina State University RED Talk, NCSU Data Science Initiative, November 6, 2019
North Carolina Chapter of the American Statistical Association, Webinar, November 1, 2019
Duke University, Langford Lecture, October 31, 2019
North Carolina State University Statistics Department Seminar, October 4, 2019
North Carolina State University ECE Interdisciplinary Distinguished Seminar Series, September 27, 2019
New York University, PRIISM Seminar, September 25, 2019
University of North Carolina at Chapel Hill, Biostatistics Seminar, August 29, 2019
SAMSI Seminar, August 28, 2019
Federal Judicial Center and Duke Law's Bolch Judicial Institute, Duke University, Workshop on Law and Technology for Judges, May 30, 2019
Yale University, MacMillan-CSAP Workshop on Quantitative Research Methods, April 25, 2019
University of Pennsylvania, Mahoney Institute for Neurosciences Seminar, April 8, 2019
Microsoft Research New England Colloquium Series, December 12, 2018
Boston University, Distinguished Speaker Series, Artificial Intelligence Research Initiative, December 10, 2018
Princeton University, Quantitative Social Science Seminar, November 16, 2018
University of North Carolina Statistics and Operations Research Department Colloquium, October 15, 2018
Research Triangle Analysts, lecture October 16, 2018
NSF Webinar: Statistics at a Crossroads: Challenges and Opportunities in the Data Science Era, October 2, 2018
University of North Carolina Health Care and NC Women in Machine Learning and Data Science MeetUp, June 2, 2018
University of North Carolina, Chapel Hill, Causal Inference Research Group Seminar, April 6, 2018
North Carolina State University, Bioinformatics Seminar, March 29, 2018
University of Maryland, Distinguished Speaker Series, Computer Science Department, December 1, 2017
Temple University, Fox School of Business, Statistics Department Colloquium, October 6, 2017
Duke University Algorithms Seminar, September 28, 2017
University of Toronto Law School, Law and Economics Colloquium, September 26, 2017
Microsoft Research, New York City, Sept 19, 2017
University at Buffalo, CSE, Distinguished Speaker Series, May 11, 2017
Columbia University, IEOR-DRO Seminar, May 9, 2017
Statistics C. V. Starr Lectureship Series, Biostatistics Department, Brown University, April 24 2017
Computational Social Science and Public Policy Colloquium, Harris School of Public Policy, University of Chicago, April 14, 2017
Statistics Department Seminar, University of Chicago, April 11, 2017
Duke Network Analysis Center Seminar, Duke University, February 6, 2017

Applied Mathematics Seminar, Duke University, October 3, 2016
Statistics and Complex Systems Seminar, University of Michigan, March 5, 2016
One Day University, New York, April 29, 2016
Urban Social Processes Workshop and Quantitative Methods Workshop, Harvard University, March 10, 2016
Center for Statistics and Machine Learning Seminar Series, Princeton University, October 20, 2015
Applied Statistics Workshop, Harvard University, October 14, 2015.
Brainstorming Session on the Next Generation of Search Engines, Berkman Center for Internet and Society, Harvard University, September 22, 2015
Machine Learning Seminar, Gatsby Unit, University College London. September 16, 2015
MIT Conversation Series, Accenture. July 10, 2015
ENAR webinar, (Eastern North American Region, International Biometric Society), May 8, 2015
Oracle Labs, Research Seminar, April 29, 2015
NYU-Poly Center for Urban Science and Progress, Research Seminar, March 12, 2015
American Express (NYC), Decision Science Seminar, March 11, 2015
Columbia University, IEOR-DRO Distinguished Seminar Series, March 10, 2015
University of Washington, Statistics Seminar, March 2, 2015
Duke University, Machine Learning Seminar, February 23, 2015
Harvard Business School, Technology and Operations Management Seminar, December 19, 2014
NYU Stern School, Department of Information, Operations & Management Sciences, IOMS Colloquium Series, December 3, 2014
UC Berkeley, Seminar in Computer Science Department, November 12, 2014
MIT Lincoln Labs, Seminar, November 4, 2014.
Brown University, Pattern Theory Group Seminar, October 22, 2014
University of Washington, Data Science Seminar, October 15, 2014
University of Alberta, Operations and Information Systems Seminar, October 3, 2014.
Carnegie Mellon University, ECE Seminar / Machine Learning Special Seminar, September 18, 2014
Harvard University, Computer Science Seminar, September 3, 2014
IBM TJ Watson Research Center, KDD Speaker Day, invited talk, August 28, 2014
MIT CSAIL CAP Meeting, May 30, 2014
MIT Lincoln Labs, Seminar, May 27, 2014
Stanford University, Operations Management Seminar, May 6, 2014
Stanford University, Institute for Research in the Social Sciences, Data Science and Inference Seminar, May 5, 2014
UMass Amherst, Machine Learning and Friends Lunch Seminar, April 29, 2014
MIT International Liaison Program Conference, Plenary Speaker, April 23, 2014
Schlumberger-Doll Research Center Seminar, April 14, 2014
Cornell University Operations Research and Industrial Engineering Seminar, April 8, 2014
Liberty Mutual Research Center Seminar, March 14, 2014
University of Pennsylvania Criminology Seminar, March 7, 2014
MIT Theory of Computation Seminar, March 4, 2014
Harvard Applied Statistics Workshop, October 9, 2013
Harvard/MIT Econometrics Workshop, September 12, 2013
MIT Lincoln Laboratory, Seminar, June 4, 2013
Massachusetts General Hospital (MGH), Quantitative Medicine Seminar, April 29, 2013
UT Austin McCombs School of Business, Research Seminar, April 5, 2013
Laboratory for Information Decision Systems (LIDS) lunchtime seminar, February 26, 2013
North Carolina State University, Statistics Department Seminar, February 28, 2013
Yale Statistics Department Seminar, January 14, 2013
Harvard High Dimensional and Correlated Data Seminar, December 17, 2012
MIT Operations Research Center Seminar, December 11, 2012
Yale Computer Science Department Seminar, December 6, 2012
Columbia University, Statistics Department Seminar, October 22, 2012
Robert H. Smith School of Business, University of Maryland, DO&IT Seminar Series, October 12, 2012
MIT Center for Collective Intelligence, July 17, 2012
Notre Dame Computer Science Department Seminar, November 3, 2011
Rutgers Statistics Department Seminar, October 26, 2011
Wharton Statistics Department Seminar, University of Pennsylvania, September 21, 2011

MIT Energy Initiative External Board Meeting Speaker, October 15, 2010
MIT Energy Initiative Seminar Series, October 12, 2010
Boston University, Probability and Statistics Seminar, October 7, 2010
Microsoft Research New England, Machine Learning Seminar, October 4, 2010
Tufts, Computer Science Seminar, September 30, 2010
ABB (Asea Brown Boveri Ltd) Corporate Research Center - United States, Lunch time Seminar, September 1, 2010.
Harvard Statistics Colloquium, April 5, 2010
MIT Operations Research Center Seminar, February 11, 2010
MIT Imaging Seminar, October 22, 2009
University of Chicago, Statistics Department Seminar, February 23, 2009
Ohio State University, Computer Science Department Seminar, February 19, 2009
Brown University, Applied Mathematics Seminar, February 17, 2009
Indiana University, Computer Science Department Seminar, March 5, 2009
University of Houston, Mathematics Department Seminar, Spring 2009
Polytechnic University, Brooklyn, Computer Science Colloquium, April 30, 2007
Columbia University, Applied Math Seminar, April 3, 2007
New York University, Theory Seminar, Computer Science Department, November 9, 2005
IBM Yorktown Heights, April 5, 2005
Rensselaer Polytechnic Institute, March 8, 2005
CCR (IDA Center for Communications Research), Princeton NJ, March 2, 2005
Institute for Advanced Study, Computer Science/Discrete Math Seminar, Princeton, February 14, 2005
Columbia University CCLS Center, February 4, 2005
Google Labs Inc., Research Seminar, October 29, 2004
New York University, Harmonic Analysis and Signal Processing Seminar, October 20, 2004
SUNY at Buffalo, Math Department Seminar, May 2004
NYU, Workshop on Computational and Biological Learning, January 16, 2004

Teaching and Mentoring

- **CS 671 / STA 671 Machine Learning, Duke**, graduate and undergraduate machine learning, 2016 (Fall), 2018 (Spring), 2019 (Spring), 2019 (Fall), 2020 (Fall), 2021 (Fall)
- **CS 290 / CS 474, Data Science Competition, Duke**, undergraduate course, 2018 (Spring), 2020 (Spring), 2021 (Spring), 2022 (Spring)
- **ME 555 Applications in Data and Materials Science**, (co-taught) 2021 (Spring)
- **Machine Learning Summer School, Duke**, 2018
- **Microsoft-DAT203x, Data Science and Machine Learning Essentials**, co-taught with Stephen Elston, free online course, edX, 2015. Over 17,500 students registered.
- **Microsoft-DAT203.2x Principals of Machine Learning**, co-taught with Stephen Elston, free online course, edX, 2016. Over 14,500 students registered.
- **Microsoft-DAT203.3x, Applied Machine Learning**, co-taught with Stephen Elston, free online course, edX, 2016. Over 8,000 students registered.
- **15.060 Data Models and Decisions, MIT**, MBA course (core course), Fall 2014, Instructor
- **15.075 Statistical Thinking and Data Analysis, MIT**, undergraduate course, Fall 2009, Fall 2010, Fall 2011, Spring 2013, Instructor.
- **15.097 Prediction: Machine Learning and Statistics, MIT**, graduate course, Spring 2012, Instructor. Course materials available on MIT Open Courseware.
- **15.060 Data Models and Decisions, MIT** MBA course (core course), Fall 2012, Instructor

- **15.064 Probability and Statistics, MIT**, Summer 2010, Summer 2011. masters student course (Leaders for Global Operations Program), Co-Instructor, 2010 and 2011
- **COMS 4771 Machine Learning, Computer Science Department, Columbia University**, Spring 2008, lectures on regression, boosting, logistic regression, and ranking.
- **Math 103 Calculus, Princeton**, Fall 2002, Fall 2001, Instructor
  - My lectures were videotaped and placed online. I was the first instructor at Princeton in the sciences to have their lectures videotaped. Class average was over 10 percentage points higher than the average of the other sections on a shared final exam that was worth 50% of their grade; this class was the top scoring class, and it scored 5 percentage points above the second highest class.
- **Math 199 Math Alive, Princeton**, Fall 2003, Teaching Assistant, responsible for the cryptography section, taught by Dr. Ingrid Daubechies
- **Wavelets Course, Program for Women in Mathematics, Institute for Advanced Study**, Summer 2002, Teaching Assistant, taught by Dr. Ingrid Daubechies
- **Physics Classes, Buffalo Seminary Women's High School**, Substitute Teacher, part-time during winter and spring, 1999, taught physics classes daily to freshmen (conceptual physics) and seniors (physics and advanced physics).

Service to Duke

Founding faculty member, AI/Materials (aiM) graduate program at Duke 2020-present, graduate admissions committee, 2021
Duke CS Department Graduate Affairs Committee, 2020-2022
Faculty Search Committees (CS, ECE, and Biostatistics), 2019-2022
Tenure Committee for Prof. David Page, Duke Biostatistics and Bioinformatics, 2020
CS Department Strategic Planning Committee, 2019
Tenure Committee for Prof. Kirsten Wickelgren, Duke Math, 2019
Graduate admissions committee, Duke CS, 2019, 2020, 2021
Working group member for white paper "Current State and Near-Term Priorities for AI-Enabled Diagnostic Support Software in Health Care," Duke-Margolis Center for Health Policy
Lead organizer for the Machine Learning seminar, 2016-present
Lead organizer of Triangle Machine Learning Day, 2018, 2019
Chair of faculty search committee, Duke CS/ECE 2017-2018
Bass connections reviewing, 2016
Grad student admissions reviewing 2015-present
Committee for a Prof. Katherine Heller's reappointment 2017-2018
Committee of Guillermo Sapiro, Vince Conitzer, and I, appointed by Provost Kornbluth to write "Computing For Humanity," 2017
Committee of 7 CS/ECE faculty members led by Carlo Tomasi to write a document similar to the above on AI, 2018
Review committee for a dean's reappointment, 2018
CS graduate awards committee, 2018
Reviewer for Data+ proposals, 2016-present
Member of numerous RIP, prelim, PhD, MS thesis, and undergraduate honors thesis committees, 2017-2018
Outreach such as hosting Duke Conversations, giving keynote for FEMMES at Duke, meeting with Visiting Committee, etc., 2018-present

Supervision

Postdocs

Dr. Aaron Fisher, co-advised with Francesca Dominici, Harvard, 2016-2019.
Dr. Keivan Sadeghzadeh, MIT Sloan, 2016.
Dr. Berk Ustun, Harvard CS, 2017-2020.

Dr. Noor-E-Alam, MIT Sloan, 2014-2015. (Now assistant professor at Northeastern University)
Dr. Ramin Moghaddass, MIT Sloan, 2013-2015 (Now assistant professor at University of Miami).
Dr. Şeyda Ertekin, MIT Sloan, 2010-2014. (Now assistant professor at Middle East Technical University).

Graduate Students

PhD student Kentaro Hoffman, UNC PhD student, 2019-2022.
MS student Pranay Jain, Duke CS MS student, 2021-present.
PhD student Rui Zhang, Duke CS PhD student, 2021-present.
PhD student Stephen Hahn, Duke ECE PhD student, 2021-present.
PhD student Jiachang Liu, Duke ECE PhD student, 2021-present.
MS student Zhicheng (Stark) Guo, Duke MS CS student, 2021-present.
MS student Xian (Jesse) Sun, Duke MS ECE student, 2020-2021.
MS student Vaishali Jain, Duke MS CS student, 2020-2021.
MS student Bin Han, Duke MS Statistics student, 2019-2020 (now PhD student at University of Washington)
MS student Neha Gupta, Duke MS Economics Computation student and then Duke Economics PhD student, 2019-present
PhD student Yingfan Wang (advised as UG, and then Duke PhD student), 2019-present
PhD Student Vittorio Orlandi, Duke Stats student, 2019-present.
PhD Student Haiyang Huang, Duke CS student, 2019-present.
MS student Henry Yuren Zhang, Duke MS Statistics student 2019-2020.
MS/PhD student Chudi Zhong, Duke MS Statistics student and then Duke CS PhD student, 2019-present
MS student Jiali Xing, Duke Economics and CS student, 2019-2020.
MS student Matias Benitez Sr., Duke Economics and CS student, 2018-2019.
MS student Chunxiao Li, Duke MS Statistics student, 2018-2019.
MS student Weiyu Yan, Duke ECE student, 2018-2019.
PhD student Usaid Awan, Duke Economics PhD student, 2018-2020.
PhD student Zhi Chen, Duke CS PhD student, 2018-present.
PhD student Jiayun Dong, Duke Economics PhD student, 2018-2019.
MS student Kangcheng Lin, Duke MS Statistics student, 2018-2019 (now at UIUC PhD program)
MS student Yang Bao, Duke Statistics student, 2018-2019.
MS student Sijia Wang, Duke ECE student, 2018-2019.
MS student Lei Chen, Duke ECE student, 2018-2020.
MS student Xiyang Hu, Duke Statistics student, 2018-2019 (now at CMU PhD program)
PhD Student Alina Barnett, Duke CS student, 2017-present.
PhD Student Harsh Parikh, Duke CS student, 2018-present.
MS Student Yameng Liu, Duke Computer Science student, 2017-2019.
PhD Student Lesia Semenova, Duke CS student, 2016-present.
PhD Student Chaofan Chen, Duke CS student, 2016-2020 (now faculty at University of Maine)
PhD Student Tianyu Wang, Duke CS student, 2016-2021 (now faculty at Fudan University)
MS Student Beau Coker, Duke Statistics student, 2017-2018 (now at Harvard PhD program)
PhD Student Marco Morucci, Duke Political Science student, 2017-2021 (now postdoc at NYU)
PhD Student Hongyu Yang, MIT EECS student, 2014-2019.
MS Student Peter Alexander Lee, MIT ORC student, 2015-2016.
MS Student Prashan Wanigasekara, MIT EECS student, 2014-2016.
MS Student Christopher Choo, Engineering and Management, 2014-2015 (now at SUTD and Singapore Grand Prix)
PhD Student Vikas Garg, MIT EECS student, 2014-2016 (co-advised with Tommi Jaakola)
PhD Student Fulton Wang, MIT EECS student, 2013-2018. (Now at Sandia National Labs)
PhD Student Berk Ustun, MIT EECS student, 2012-2017 (Now faculty at UCSD)
PhD Student Stefano Tracà, MIT ORC student, 2012-2018 (now working at Disney Research)

PhD Student Siong Thye Goh, MIT ORC student, 2012-2018.

PhD Student Tong Wang, MIT EECS student, 2012-2016. (Now faculty at University of Iowa)

Project Student Ashia Wilson, MIT Sloan, 2012. (5 months before starting a PhD program at Berkeley)

PhD Student Theja Tulabandhula, MIT EECS student, 2010-2014. (Now senior lecturer at University of Sydney Business School)

PhD Student Ben Letham, MIT ORC student, 2010-2015. (Now at Facebook)

PhD Student Allison Chang, MIT ORC student, co-supervised with Dimitris Bertsimas, 2009-2012 (now at MIT Lincoln Labs).

Masters Student William Harris, MIT ORC Student, co-advised with Michael Ricard, 2014-2015 (now in the US military)

MS Student Oscar Moll, MIT CSAIL student, 2010-2011.

MS Graduate Research Assistant, Nandini Bhardwaj, Columbia & Con Edison Secondary Events Project, 2008.

Masters Project Course, Jawwad Sultan, Columbia & Con Edison Secondary Events Project, Fall 2007.

Summer Students, Supervision of 2 masters students and 1 undergraduate. Columbia & Con Edison Secondary Events Project, Summer 2007.

Undergraduate Students

Duke undergraduate, Jessie Ou, 2021.

Duke undergraduate, Jerry Fang, 2021.

Duke undergraduate, Vaibhav Sharma, 2021.

Duke undergraduate, Rui Xin, 2021.

Duke undergraduate, Harsha Srijay, 2021.

Duke undergraduate, Vijit Singh, 2021.

Duke undergraduate, Caleb Kornfeld, 2021.

Duke undergraduate, George Wang, 2021.

Duke undergraduate, Alexander Oesterling, 2021.

Duke undergraduate, Haoning Jiang, 2021.

Duke undergraduate, Lily Zhu, 2021.

Duke undergraduate, Yunyao Zhu, 2021.

Duke undergraduate, Jerry Liu, 2020-2021.

Duke undergraduate, Nathan O'Hara, 2020.

Duke undergraduate, Krystal Hu, 2020.

Duke undergraduate, Angikar Ghosal, 2020-2021.

Duke undergraduate, Thomas Howell, 2020-2021.

Duke undergraduate, Edwin Agnew, 2020-2021.

Duke undergraduate, Benjamin Burnette, 2020-2021.

Duke undergraduate, Jordan Diamond, 2020-2021.

Duke undergraduate, Reed Chen, 2020.

Duke undergraduate, Kari Larson, 2020-2021.

Duke undergraduate, Brandon Zhao, 2019-2021.

Duke undergraduate, Alexander Rubin, 2019.

Duke undergraduate, Feroze Mohideen, 2019.

Duke undergraduate, Diane Hu, 2019-2020.

Duke undergraduate, Isaac Zhang, 2019-2021.

Duke undergraduate, Andre Wang, 2019-2021.

Duke undergraduate, Bhrij Patel, 2019-2021.

Duke undergraduate, Jake Shulman, 2019.

Duke undergraduate, Kenny Green, 2019.

Duke undergraduate, Jerry Pan, 2018-2019.

Duke undergraduate, Alexandru Damian, 2018-2020.
Duke undergraduate, Nikhil Ravi, 2018-2020.
Duke undergraduate, Sachit Menon, 2018-2020.
Duke undergraduate, Chris Suh, 2018-2019.
Duke undergraduate, McCourt Hu, 2018-2019.
Duke undergraduate, Webster Bei, 2018-2020.
Duke undergraduate, Jerry Chia Rui Chang, 2018.
Duke undergraduate, Wilson Zhang, 2018-2019.
Duke undergraduate, Divya Koyyalagunta, 2018-2019.
Duke undergraduate, Anna Sun, 2018-present.
Duke undergraduate, Peter Hase, 2018-2019.
Duke undergraduate, Daniel Tau, 2017-2019.
Duke undergraduate, Caroline Wang, 2017-2020.
Duke undergraduate, Hao Liu, summer 2017.
Duke undergraduate, Oscar Li, 2017-2019.
MIT undergraduate, Chelsea Ge, summer 2014.
MIT undergraduate, Jeffrey Chan, spring-fall 2014.
MIT undergraduate, Jiaming Zeng, fall-spring 2015.
MIT undergraduate, Shawn Qian, summer-fall 2012.
Undergraduate exchange student, Yining Wang, spring 2013.
PhD student at MIT and previously undergraduate from Arizona State University, Lydia Letham, summer 2012, summer 2014.
MIT undergraduate project courses, three students (Kang Zhang, Arash Delijani, Kevin Pang) 2011-2012.
Undergraduate Visiting Student from Ecole Centrale Paris (through MISTI), Fabrice Vegetti, 2012.
Undergraduate Visiting Student from Ecole Centrale Paris (through MISTI), Adel Basli, 2011.
Undergraduate project course on Collaborative Filtering, Association Rules and Information Retrieval, Eugene Kogan, Columbia University, co-supervision with Dr. Ansaf Salleb-Aouissi, Spring 2008.
Undergraduate thesis advisement at Princeton, Krysta Svore, entitled “Multiscale Image Processing Using Single and Double Gaussian Techniques, and Hidden Markov Models,” 2001-2002.

Thesis/Prelim Committees (not including my students)

Qual committee for Duke MEMS PhD student, Jake Peloquin, 2022
Qual committee for Duke ECE PhD student, Lin Duan, 2021
Prelim committee for Duke CS student Qinwen Huang, 2021
PhD defense committee for Duke ECE student Haibei Zhu, 2021
Prelim committee for Duke ECE student Mojtaba Zarei, 2021
Undergraduate honors thesis committee for Yunyao Zhu, 2021
Prelim committee for Duke MEMS student Peiyi Chen, 2021
Prelim committee for Duke Biomedical Engineering and Radiology PhD student Yinhao Ren, 2021
Prelim & PhD thesis committee for Duke ECE PhD student Bohao Huang, 2019, 2020
Prelim committee for Duke MEMS PhD student Bingyin Hu, 2020
Prelim committee for Duke ECE PhD student Jiachang Liu, 2019
Qual committee for Duke ECE PhD student Kavinayan Sivakumar, 2019
Prelim committee for Duke Statistics PhD student Filipe Ossa, 2019
Prelim & PhD Thesis committee for Duke Economics Student Usaid Awan, 2019, 2022
RIP committee for Duke ECE PhD student Jerry Wang, 2019
Prelim committee for Duke Fuqua PhD student Shuyu Chen, 2019
RIP & Prelim committee for Duke CS PhD student Shuai Yuan, 2019, 2020
Qual & Prelim committee for Duke ECE PhD student Yuting Ng, 2019, 2021
Qual committee for Duke ECE PhD student Qian Huang, 2019
Qual, Prelim, & PhD committees for Duke ECE PhD student Ghassen Jerfel, 2019, 2020, 2021

Prelim & PhD thesis committee for Duke ECE PhD student Wanyi Fu, 2019, 2021
Qualifying committee for Duke ECE PhD student Claire Lin, 2019
Prelim committee for Duke CS PhD student Shuzhi Yu, 2018
Prelim committee for Duke CS PhD student Swarna Ravindran, 2018
Prelim committee for PhD student Paidamoyo Chapfuwa, ECE PhD student, 2018
RIP committee and prelim committee for Duke CS PhD student Andrew Lee, 2017, 2018
Graduation with honors committee for Duke undergraduate Peter Hase, 2018
Graduation with honors committee for Duke undergraduate Wuming Zhang, 2018
Graduation with honors committee for Duke undergraduate Tianlin Duan, 2018
RIP for Duke PhD student Xiaonan Hu, 2017, 2018
Prelim committee for Duke PhD student Greg Spell, 2018
RIP, Prelim and Thesis committees for Duke PhD student Rachel Draelos, 2017, 2019, 2021
Prelim and thesis committees for Duke PhD student Stavros Sintos, 2017, 2020
RIP and Prelim committee for Duke CS PhD student Zilong Tan, PhD in fall 2018
Thesis committee for Stanford CS student Himabindu Lakkaraju, PhD in spring 2018
RIP committee for Duke CS PhD student Xiaonan Hu, 2017
RIP committee for Duke CS PhD student Shuzhi Yu, 2017
RIP, prelim, PhD committees for Duke CS PhD student Abe Frandsen, 2018, 2019, 2022
Thesis committee for Duke CS undergraduate with distinction Aditya Mukund, 2017
Thesis committee for Duke CS masters student Guan-Wun Hao, 2017
Thesis committee for Duke CS masters student Mona Prakash, 2016
Thesis committee for Duke statistics masters student Emily Shao, 2017
Thesis committee for Duke statistics masters student Sanjay Harihanan, 2017
Thesis committee for Duke ECE Phd student Jordan Hashemi, 2017
Thesis committee for Duke ECE Phd student Zhuoqing Chang, 2017, 2018, 2020
Thesis reader for Harvard CS undergraduate Nicholas Larus-Stone, 2017
Thesis committee for Duke PhD student Shan Shan, 2017
Thesis committee for Duke PhD student Narayanan Rengaswamy, 2017.
Thesis committee for MIT PhD student Yingxiang Yang, 2015.
Thesis committee for MIT PhD student Been Kim, PhD in spring 2015.
Thesis committee for MIT PhD student Anima Singh, PhD in spring 2015.
Thesis committee for Pannaga Shivaswamy at Columbia University CS, PhD in spring 2009.

Society Memberships

-INFORMS
-International Machine Learning Society
-American Statistical Association (ASA)
-Institute of Mathematical Statistics (IMS)
-Association for the Advancement of Artificial Intelligence (AAAI)
-Association for Computing Machinery (ACM)
-American Association for the Advancement of Science (AAAS)
-Society for Causal Inference (SCI)